arXiv:1505.02294v3 [stat.ML] 30 Nov 2015

Estimation with Norm Regularization

Arindam Banerjee Sheng Chen Farideh Fazayeli Vidyashankar Sivakumar

{banerjee,shengc,farideh,sivakuma@cs.umn.edu}

Department of Computer Science & Engineering
University of Minnesota, Twin Cities

December 1, 2015

Abstract

Analysis of non-asymptotic estimation error and structured statistical recovery based on norm regularized regression, such as Lasso, needs to consider four aspects: the norm, the loss function, the design matrix, and the noise model. This paper presents generalizations of such estimation error analysis on all four aspects compared to the existing literature. We characterize the restricted error set where the estimation error vector lies, establish relations between error sets for the constrained and regularized problems, and present an estimation error bound applicable to *any* norm. Precise characterizations of the bound is presented for isotropic as well as anisotropic subGaussian design matrices, subGaussian noise models, and convex loss functions, including least squares and generalized linear models. Generic chaining and associated results play an important role in the analysis. A key result from the analysis is that the sample complexity of all such estimators depends on the Gaussian width of a spherical cap corresponding to the restricted error set. Further, once the number of samples n crosses the required sample complexity, the estimation error decreases as $\frac{c}{\sqrt{n}}$, where c depends on the Gaussian width of the unit norm ball.

1 Introduction

Over the past decade, progress has been made in developing non-asymptotic bounds on the estimation error of structured parameters based on norm regularized regression. Such estimators are usually of the form [39, 29, 9]:

$$\hat{\theta}_{\lambda_n} \;=\; \operatorname*{argmin}_{\theta\in\mathbb{R}^p} \mathcal{L}(\theta; Z^n) + \lambda_n R(\theta)\;, \tag{1}$$

where $R(\theta)$ is a suitable norm, $\mathcal{L}(\cdot)$ is a suitable loss function, $Z^n = \{(y_i, X_i)\}_{i=1}^n$ where $y_i \in \mathbb{R}, X_i \in \mathbb{R}^p$ is the training set, and $\lambda_n > 0$ is a regularization parameter. The optimal parameter θ^* is often assumed to be 'structured,' usually characterized or approximated as a small value according to some norm $R(\cdot)$. Recent work has viewed such characterizations in terms of atomic norms, which give the tightest convex relaxation of a structured set of atoms in which θ^* belongs [14]. Since $\hat{\theta}_{\lambda_n}$ is an estimate of the optimal structure θ^*, the focus has been on bounding a suitable measure of the error vector $\hat{\Delta}_n = (\hat{\theta}_{\lambda_n} - \theta^*)$, e.g., the L_2 norm $\|\hat{\Delta}_n\|_2$.

To understand the state-of-the-art on non-asymptotic bounds on the estimation error for norm-regularized regression, four aspects of (1) need to be considered: (i) the norm $R(\theta)$, (ii) properties of the design matrix $X = [X_1 \cdots X_n]^T \in \mathbb{R}^{n\times p}$, (iii) the loss function $\mathcal{L}(\cdot)$, and (iv) the noise model, typically in terms of $\omega_i = y_i - E[y|X_i]$. Most of the literature has focused on a linear model: $y = X\theta + \omega$, and a squared-loss function: $\mathcal{L}(\theta; Z^n) = \frac{1}{n}\|y - X\theta\|_2^2 = \frac{1}{n}\sum_{i=1}^n (y_i - \langle\theta, X_i\rangle)^2$. Early work on such estimators focussed on

the L_1 norm [45, 43, 26], and led to sufficient conditions on the design matrix X, including the restricted-isometry properties (RIP) [13, 12] and restricted eigenvalue (RE) conditions [6, 29, 33]. While much of the development has focussed on isotropic Gaussian design matrices, recent work has extended the analysis for L_1 norm to correlated Gaussian designs [33] as well as anisotropic sub-Gaussian design matrices [34].

Building on such development, [29] presents a unified framework for the case of decomposable norms and also considers generalized linear models (GLMs) for certain norms such as L_1. Two key insights are offered in [29]: first, the error vector $\hat{\Delta}_n$ lies in a restricted set, a cone or a star, for suitably large λ_n, and second, the loss function needs to satisfy restricted strong convexity (RSC), a generalization of the RE condition, on the restricted error set for the analysis to work out.

For isotropic Gaussian design matrices, additional progress has been made. [14] considers a constrained estimation formulation for all atomic norms, where the gain condition, equivalent to the RE condition, uses Gordon's inequality [19, 20, 24] and is succinctly represented in terms of the Gaussian width of the intersection of the cone of the error set and a unit ball/sphere. [31] considers three related formulations for generalized Lasso problems, establish recovery guarantees based on Gordon's inequality, and quantities related to the Gaussian width. Sharper analysis for recovery has been considered in [1], yielding a precise characterization of phase transition behavior using quantities related to the Gaussian width. [32] consider a linear programming estimator in a 1-bit compressed sensing setting and, interestingly, the concept of Gaussian width shows up in the analysis. In spite of the advances, with a few notable exceptions [40, 42], most existing results are restricted to isotropic Gaussian design matrices. Further, while a suitable scale for λ_n is known for special cases such as the L_1, a general analysis applicable to any norm $R(\cdot)$ has not been explored in the literature.

In this paper, we consider structured estimation problems with norm regularization of the form (1), and present a unified analysis which substantially generalizes existing results on all four pertinent aspects: the norm, the design matrix, the loss, and the noise model. The analysis we present applies to *all* norms, and the results can be divided into three groups: characterization of the error set and recovery guarantees, characterization of the regularization parameter λ_n, and characterization of the restricted eigenvalue conditions or restricted strong convexity. We provide a summary of the key results below.

Restricted error set: We start with a characterization of the error set E_r in which the error vector $\hat{\Delta}_n$ belongs. For a suitably large λ_n, we show that $\hat{\Delta}_n$ belongs to the restricted error set

$$E_r = \left\{ \Delta \in \mathbb{R}^p \;\middle|\; R(\theta^* + \Delta) \leq R(\theta^*) + \frac{1}{\beta} R(\Delta) \right\} , \tag{2}$$

where $\beta > 1$ is a constant. The restricted error set has interesting structure, and forms the basis of subsequent analysis for bounds on $\|\hat{\Delta}_n\|_2$.

Regularized vs. constrained estimators: As an alternative to regularized estimators, the literature has considered constrained estimators which directly focus on minimizing $R(\theta)$ under suitable constraints determined by the noise $(y - X\theta)$ and/or the design matrix X [13, 6, 14, 15]. A recent example of such a constrained estimator is the generalized Dantzig selector (GDS) [15], which generalizes the Dantzig selector [11] corresponding to the L_1 norm, and is given by:

$$\hat{\theta}_{\gamma_n} \;=\; \underset{\theta \in \mathbb{R}^p}{\operatorname{argmin}} \, R(\theta) \quad \text{s.t.} \quad R^*(X^T(y - X\theta^*)) \leq \gamma_n \, , \tag{3}$$

where $R^*(\cdot)$ denotes the dual norm of $R(\cdot)$. One can show [14, 15] that the restricted error set for such constrained estimators are of the form:

$$E_c = \{ \Delta \in \mathbb{R}^p \mid R(\theta^* + \Delta) \leq R(\theta^*) \} \, . \tag{4}$$

One can readily see that E_r is larger than E_c, i.e., $E_c \subseteq E_r$, and E_r approaches E_c as β increases. We establish a geometric relationship between the two sets, which will possibly help in transforming analysis done on regularized estimators as in (1) to corresponding constrained estimators as in (3) and vice versa. Let B_2^p be a L_2 ball of radius 1 in $\mathbb{R}^p$. Then, with $A_r = E_r \cap B_2^p$, $A_c = E_c \cap B_2^p$, and $\bar{A}_c = \text{cone}(E_c) \cap B_2^p$, assuming $\|\theta^*\|_2 = 1, \beta = 2$, we show that

$$w(A_c) \leq w(A_r) \leq 3w(\bar{A}_c) \ , \tag{5}$$

where $w(A) = E_g[\sup_{a \in A} \langle a, g \rangle]$, with $g \sim N(0, \mathbb{I}_{p \times p})$ being an isotropic Gaussian vector, denotes the Gaussian width[1] of the set A [24, 5, 14, 37]. Note that A_r corresponds to the spherical cap of the error set E_r, and A_c corresponds to the spherical cap of the error cone $\text{cone}(E_c)$ at the unit ball. Interestingly, the above relationship between the widths of these spherical caps is geometric, and applies for any norm $R(\cdot)$.

We establish a more general version of the above relationship. Let ρB_2^p denote a L_2 ball of any radius ρ in $\mathbb{R}^p$. Then, with $A_r^{(\rho)} = E_r \cap \rho B_2^p$, $A_c^{(\rho)} = E_c \cap \rho B_2^p$, and $\bar{A}_c^{(\rho)} = \text{cone}(E_c) \cap \rho B_2^p$, we show that

$$w(A_c^{(\rho)}) \leq w(A_r^{(\rho)}) \leq \left(1 + \frac{2}{\beta - 1} \frac{\|\theta^*\|_2}{\rho}\right) w(\bar{A}_c^{(\rho)}) \ . \tag{6}$$

As before, except for the scaling constants, the relationship between the restricted error sets is geometric, and does not change based on the choice of the norm $R(\cdot)$.

For the special case of L_1 norm, [6] considered a simultaneous analysis of the Lasso and the Dantzig selector, and characterized the structure of the error sets for regularized and constrained sets for the special case of L_1 norm. Further, while the characterization in [6] was also geometric, it was not based on Gaussian widths. In contrast, our results apply to any norm, not just L_1, and the geometric characterization is based on Gaussian widths. The utility of the Gaussian width based characterization becomes evident later when we establish sample complexity results for Gaussian and sub-Gaussian random matrices in terms of Gaussian widths of spherical caps.

Bounds on estimation error: We establish bounds on the estimation error $\hat{\Delta}_n$ under two assumptions, which are subsequently shown to hold with high probability for sub-Gaussian designs and noise models. The first assumption is that the regularization parameter λ_n is suitably large. In particular, for any $\beta > 1$, the regularization parameter λ_n needs to satisfy

$$\lambda_n \geq \beta R^*(\nabla \mathcal{L}(\theta^*; Z^n)) \ , \tag{7}$$

where $R^*(\cdot)$ denotes the dual norm of $R(\cdot)$. The second assumption is that the design matrix $X \in \mathbb{R}^{n \times p}$ satisfies the restricted strong convexity (RSC) condition [6, 29] in the error set E_r, in particular, there exists a suitable constant $\kappa > 0$ so that

$$\delta \mathcal{L}(\Delta, \theta^*) \triangleq \mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) - \langle \nabla \mathcal{L}(\theta^*), \Delta \rangle \geq \kappa \|\Delta\|_2^2 \quad \forall \Delta \in E_r \ . \tag{8}$$

With such suitably large λ_n and $\mathcal{L}$ satisfying the RSC condition, we establish the following bound:

$$\|\hat{\Delta}_n\|_2 \leq c \psi(E_r) \frac{\lambda_n}{\kappa} \ , \tag{9}$$

where $\psi(E_r) = \sup_{u \in E_r} \frac{R(u)}{\|u\|_2}$ is a norm compatibility constant [29], and $c > 0$ is a constant. Note that the above bound is deterministic, but relies on assumptions on λ_n and κ. So, we focus on characterizations of λ_n and κ which hold with high probability for sub-Gaussian design matrices X and sub-Gaussian noise ω. Recent work in [36] has extended the analyses for sub-exponential distributions.

[1] A gentle exposition to Gaussian width and some of its properties is given in Appendix A.

Bounds on the regularization parameter λ_n: From (7) above, for the analysis to work, one needs to have $\lambda_n \geq \beta R^*(\nabla \mathcal{L}(\theta^*; Z^n))$. There are a few challenges in getting a suitable bound for λ_n. First, the bound depends on θ^*, but θ^* is unknown and is the quantity one is interested in estimating. Second, the bound depends on Z^n, the samples, and is hence random. The goal will be to bound the expectation $E[R^*(\nabla \mathcal{L}(\theta^*; Z^n))]$ over all samples of size n, and obtain high-probability deviation bounds. Third, since the bound relies on the (dual) norm $R^*(\cdot)$ of a p-dimensional random vector, without proper care, the lower bound on λ_n may end up having a large scaling dependency, say $\sqrt{p}$, on the ambient dimensionality. Since the error bound in (9) is directly proportional to λ_n, such dependencies will lead to weak bounds.

In Section 3, we characterize the expectation $E[R^*(\nabla \mathcal{L}(\theta^*; Z^n))]$ in terms of the geometry of the unit norm-ball of R, which leads to a sharp bound. Let $\Omega_R = \{u \in \mathbb{R}^p | R(u) \leq 1\}$ denote the unit norm-ball. Then, for sub-Gaussian design matrices and squared loss, we show that

$$E[R^*(\nabla \mathcal{L}(\theta^*; Z^n))] \leq \frac{c}{\sqrt{n}} w(\Omega_R) \, , \tag{10}$$

which scales as the Gaussian width of Ω_R. Interestingly, for sub-Gaussian designs, one obtains the results in terms of the 'sub-Gaussian width' of the unit norm-ball, which can be upper bounded by a constant times the Gaussian width using generic chaining [37]. The result can be extended to the case of anisotropic sub-Gaussian designs, where the constant c starts depending on the maximum eigenvalue (operator norm) of the corresponding covariance matrix. Further, one can get high-probability versions of these bounds using related advances in generic chaining [37, 38]. The results can also be extended to general convex losses, such as those from generalized linear models.

The above characterization allows one to choose $\lambda_n \geq \frac{c}{\sqrt{n}} w(\Omega_R)$. For the special case of L_1 regularization, Ω_R is the unit L_1 norm ball, and the corresponding Gaussian width $w(\Omega_R) \leq c_1 \sqrt{\log p}$, which explains the $\sqrt{\log p}$ term one finds in existing bounds for Lasso [29, 9]. When working with other norms, one simply needs to get an upper bound on the corresponding $w(\Omega_R)$.

Restricted eigenvalue conditions: When the loss function under consideration is the squared loss, the RSC condition in (8) reduces to the restricted eigenvalue (RE) condition on the design matrix. Our analysis focuses on establishing the RE condition on $\bar{A}_r = \text{cone}(E_r) \cap S^{p-1}$, the spherical cap obtained by intersecting the cone of the error set with the unit hypersphere, since it implies the RE condition on E_r. For isotropic sub-Gaussian design matrices, a stronger two-sided restricted isometry property (RIP) holds, i.e., with high probability, for any $A \subseteq S^{p-1}$, we have

$$1 - c\frac{w(A)}{\sqrt{n}} \leq \inf_{u \in A} \frac{1}{n} \|Xu\|^2 \leq \sup_{u \in A} \frac{1}{n} \|Xu\|^2 \leq 1 + c\frac{w(A)}{\sqrt{n}} \tag{11}$$

where $w(A)$ is the Gaussian width of A. Thus, for say $n_0 = 4c^2 w^2(\bar{A}_r)$, and for $n > n_0$, an RE condition of the form

$$\inf_{u \in \bar{A}_r} \frac{1}{n} \|Xu\|_2^2 \;\; \geq \;\; 1/2 \, , \tag{12}$$

is satisfied with high probability. Instead of the constant to be $1/2$, one can have any constant less than 1, with suitable increase in n_0. Thus, one does not need to treat the RE condition as an assumption for isotropic sub-Gaussian designs—it always holds with high probability, with the phase transition happening at $O(w^2(\bar{A}_r))$ samples. The RIP results can be generalized to anisotropic sub-Gaussian designs, where additional constants depending on the restricted eigenvalues of the anisotropic covariance matrix Σ show up, but the form of the bound stays similar. Our analysis techniques for the RIP results are based on generic chaining [37, 38], in particular a specific form developed in [21, 28].

Generalized linear models and restricted strong convexity: For convex loss functions, such as those coming from generalized linear models (GLMs), the sample complexity and associated phase transition behavior

is determined by the Restricted Strong Convexity (RSC) condition [29]. By generalizing our argument for RE conditions corresponding to square loss, we show that the RSC conditions are going to be satisfied for convex losses for subGaussian designs at the same order of sample complexity as that for squared loss. In particular, for we show a high probability lower bound of the form

$$\inf_{u \in A} \ \delta\mathcal{L}(u, \theta^*) \geq c_1 - c_2 \frac{w(A)}{\sqrt{n}} \ , \tag{13}$$

where the constants $c_1, c_2 > 0$ depend on the tail probabilities of the design matrix distribution. Specializing the result to $\bar{A}_r = \mathrm{cone}(E_r) \cap S^{p-1}$, we note that the sample complexity still scales as $O(w^2(\bar{A}_r))$, similar to the case of squared loss. The result is thus a considerable generalization of earlier results on convex losses, such as GLMs, which had looked at specific norms and associated cones and/or did not express the results in terms of the Gaussian width of A [29].

Putting everything together: With the above results in place, from (9), the main bound takes the form

$$\|\hat{\Delta}_n\|_2 \leq c \frac{\psi(E_r)}{\left[c_1 - c_2 \frac{w(\bar{A}_r)}{\sqrt{n}}\right]_+} \frac{w(\Omega_R)}{\sqrt{n}} \tag{14}$$

with high probability, where $w(\Omega_R)$ is the Gaussian width of the unit norm ball, $w(\bar{A}_r)$ is the Gaussian width of the spherical cap corresponding to the error set $\mathrm{cone}(E_r)$, and the result is valid only when $n > n_0 = O(w^2(\bar{A}_r))$ which corresponds to the sample complexity. For the special case of L_1 norm, i.e., Lasso, the sample complexity n_0 is of the order $w^2(\bar{A}_r) = O(s \log p)$. Further, $w(\Omega_R) = \sqrt{\log p}$ and $\psi(E_r) = \sqrt{s}$. Plugging in these values, choosing $\beta = 2$, for $n > c_3 s \log p$, the bound $\|\hat{\Delta}_n\|_2 \leq c\sqrt{\frac{s \log p}{n}}$ holds with probability. For other norms, one can simply plug-in the widths to get the corresponding sample complexity and non-asymptotic error bounds.

The rest of the paper is organized as follows: Section 2 presents results on the restricted error set and deterministic error bounds under suitable bounds on the regularization parameter λ_n and RSC assumptions. Section 3 presents a characterization of λ_n in terms of the Gaussian width of the unit norm ball for Gaussian as well as sub-Gaussian designs and noise. Section 4 proves RE conditions and associated sample complexity results corresponding to squared loss functions. Results are presented for subGaussian designs, including anisotropic and correlated cases, and always in terms of the Gaussian width of the spherical cap corresponding to the error set. Section 6 presents RSC conditions corresponding to general convex losses arising from generalized linear models, and the results are again in terms of the Gaussian width of the spherical cap corresponding to the error set. We conclude in Section 7. All technical arguments and proofs are in the appendix, along with a gentle exposition to Gaussian widths and related results.

A brief word on the notation used. We denote random matrices as X, and random vectors as X_i where i may be an index to a row or column of a random matrix. Vector norms are denoted as $\|\cdot\|$, e.g., $\|X_i\|_2$ for a (random) vector X_i, and norms of random variables are denoted as $\|\!|\cdot|\!\|$, e.g., $\|\!|X|\!\|_2 = E[\|X\|_2]$.

2 Restricted Error Set and Recovery Guarantees

In this section, we give a characterization of the restricted error set E_r in which the error vector $\hat{\Delta}_n = (\hat{\theta}_{\lambda_n} - \theta^*)$ lies, establish clear relationships between the error sets for the regularized and constrained problems, and finally establish upper bounds on the estimation error. The error bound is deterministic, but has quantities which involve θ^*, X, ω, for which we develop high probability bounds in Sections 3, 4, and 6.

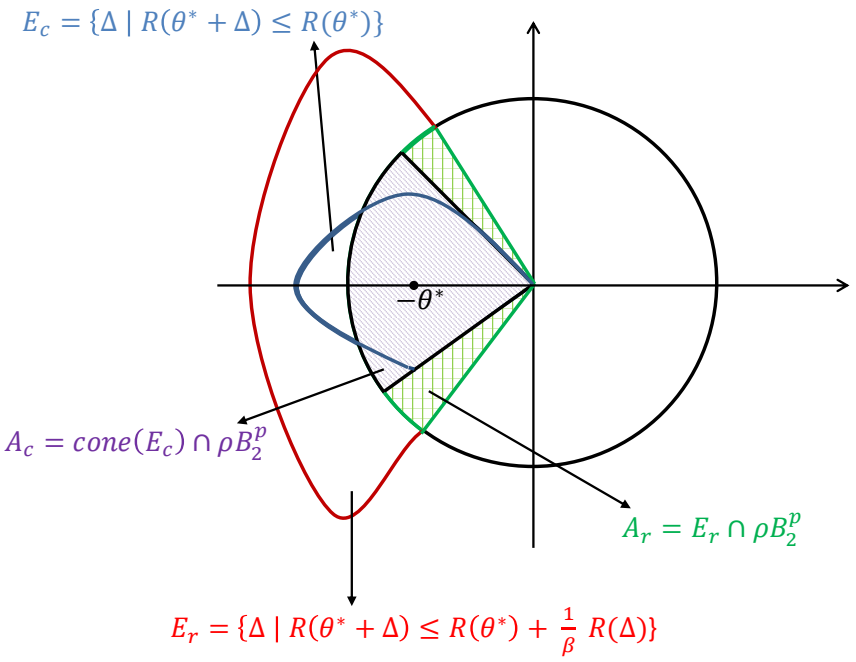

Figure 1: Relationship between error set of the regularized problem (A_r, green region) and the constrained problem (A_c, gray region) after intersection with a ball of radius ρ. While A_r will be larger in general, it will be within a constant factor of A_c in terms of Gaussian width (best viewed in color).

2.1 The Restricted Error Set and the Error Cone

We start with a characterization of the restricted error set E_r where $\hat{\Delta}_n$ will belong.

Lemma 1 *For any $\beta > 1$, assuming*

$$\lambda_n \geq \beta R^*(\nabla \mathcal{L}(\theta^*; Z^n)) \,, \tag{15}$$

where $R^*(\cdot)$ is the dual norm of $R(\cdot)$. Then the error vector $\hat{\Delta}_n = \hat{\theta}_{\lambda_n} - \theta^*$ belongs to the set

$$E_r = E_r(\theta^*, \beta) = \left\{ \Delta \in \mathbb{R}^p \;\middle|\; R(\theta^* + \Delta) \leq R(\theta^*) + \frac{1}{\beta} R(\Delta) \right\} \,. \tag{16}$$

The restricted error set E_r need not be convex for general norms. Interestingly, for $\beta = 1$, the inequality in (16) is just the triangle inequality, and is satisfied by all Δ. Note that $\beta > 1$ restricts the set of Δ which satisfy the inequality, yielding the restricted error set. In particular, Δ cannot go in the direction of θ^*, i.e., $\Delta \neq \alpha\theta^*$ for any $\alpha > 0$. Further, note that the condition in (15) is similar to that in [29] for $\beta = 2$, but the above characterization holds for any norm, not just decomposable norms [29].

While E_r need not be a convex set, we establish a relationship between E_r and the error set E_c corresponding to constrained estimators [13, 6, 14, 15]. A recent example of such a constrained estimator is the generalized Dantzig selector (GDS) [15] given by:

$$\hat{\theta}_{\gamma_n} \;=\; \underset{\theta \in \mathbb{R}^p}{\operatorname{argmin}}\; R(\theta) \quad \text{s.t.} \quad R^*(X^T(y - X\theta^*)) \leq \gamma_n \,, \tag{17}$$

where $R^*(\cdot)$ denotes the dual norm of $R(\cdot)$. One can show [14, 15] that the restricted error set for such constrained estimators [14, 15, 40] are of the form:

$$E_c = \{\Delta \in \mathbb{R}^p \mid R(\theta^* + \Delta) \leq R(\theta^*)\} \,. \tag{18}$$

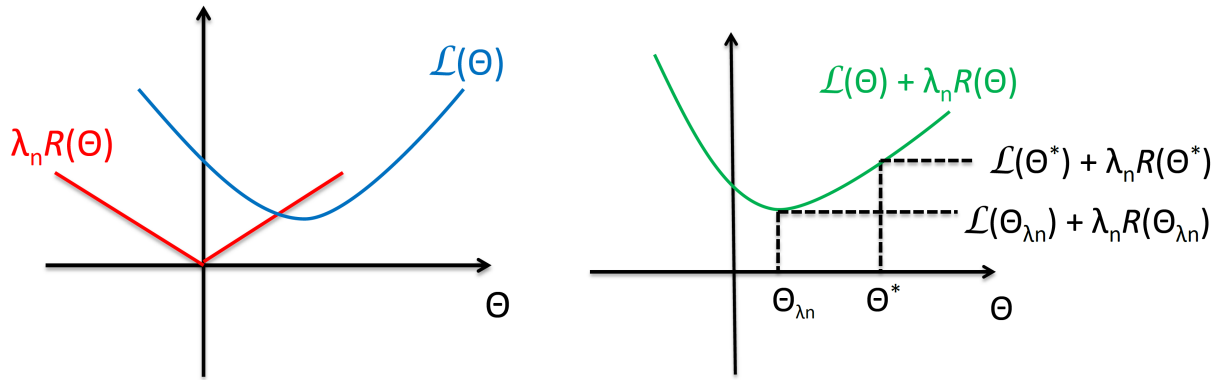

(a) Loss and regularizer. (b) Overall objective.

Figure 2: Schematic for norm regularized objective functions considered. The finite sample estimate $\hat{\theta}_{\lambda_n}$ has lower empirical loss than the optimum θ^*. Bounding the difference between the losses yields a bound on $\|\hat{\theta}_{\lambda_n} - \theta^*\|$.

By definition, it is easy to see that E_c is always convex, and that $E_c \subseteq E_r$, as shown schematically in Figure 1.

The following results establishes a relationship between E_r and E_c in terms of their Gaussian widths.

Theorem 1 *Let $A_c^{(\rho)} = E_c \cap \rho B_2^p$, $A_r^{(\rho)} = E_r \cap \rho B_2^p$, and $\bar{A}_c^{(\rho)} = C_c \cap \rho B_2^p$, where $\rho B_2^p = \{u | \|u\|_2 \leq \rho\}$ is the L_2 ball of any radius $\rho > 0$. Then, for any $\beta > 1$ we have*

$$w(A_c^{(\rho)}) \leq w(A_r^{(\rho)}) \leq \left(1 + \frac{2}{\beta - 1} \frac{\|\theta^*\|_2}{\rho}\right) w(\bar{A}_c^{(\rho)}) \,, \tag{19}$$

where $w(A)$ denotes the Gaussian width of any set A given by: $w(A) = E_g \left[\sup_{a \in A} \langle a, g \rangle\right]$, where g is an isotropic Gaussian random vector, i.e., $g \sim N(0, \mathbb{I}_{p \times p})$.

Thus, the Gaussian width of the error sets of regularized and constrained problems are closely related. See Figure 1 for more details. In particular, for $\|\theta^*\|_2 = 1$, with $\rho = 1, \beta = 2$, we have $w(A_c) \leq w(A_r) \leq 3w(\bar{A}_c)$ as introduced in Section 1. Related observations have been made for the special case of the L_1 norm [6], although past work did not provide an explicit characterization in terms of Gaussian widths. The result also suggests that it is possible to move between the error analysis of the regularized and the constrained versions of the estimation problem.

2.2 Recovery Guarantees

In order to establish recovery guarantees, we start by assuming that restricted strong convexity (RSC) is satisfied by the loss function in E_r, the error set, so that for any $\Delta \in E_r$, there exists a suitable constant κ so that

$$\delta\mathcal{L}(\Delta, \theta^*) \triangleq \mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) - \langle \nabla\mathcal{L}(\theta^*), \Delta \rangle \geq \kappa \|\Delta\|_2^2 \,. \tag{20}$$

In Sections 4 and 6, we establish precise forms of the RSC condition for a wide variety of design matrices and loss functions. In order to establish recovery guarantees, we focus on the quantity

$$\mathcal{F}(\Delta) = \mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) + \lambda_n (R(\theta^* + \Delta) - R(\theta^*)) \,. \tag{21}$$

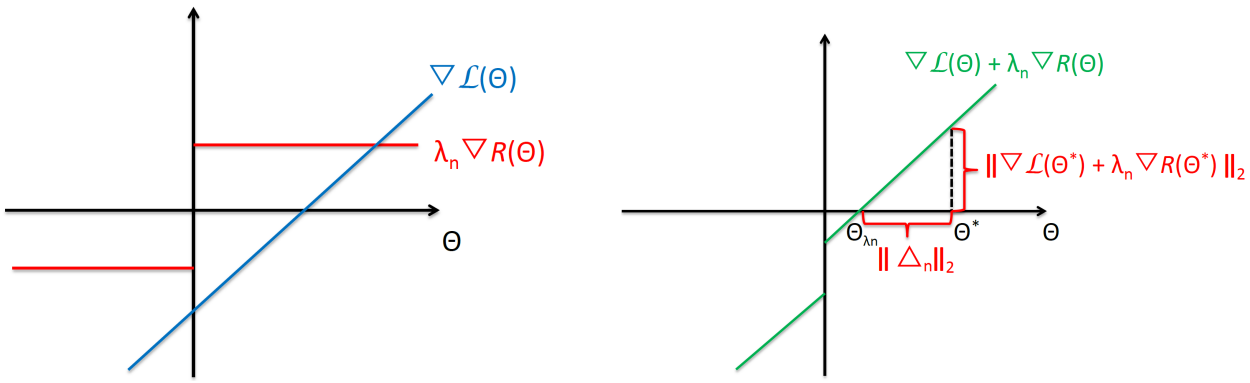

(a) Gradients of loss and norm.

(b) Gradient of the overall objective

Figure 3: Schematic illustrating the error bound in Lemma 2. Under restricted strong convexity (RSC) of the loss function in the error set E_r, the error $\|\hat{\Delta}_n\|_2$ can be bounded in terms of the gradient of the overall objective evaluated at θ^*.

Since $\hat{\theta}_{\lambda_n} = \theta^* + \hat{\Delta}_n$ is the estimated parameter, i.e., $\hat{\theta}_{\lambda_n}$ is the minimum of the objective, we clearly have $\mathcal{F}(\hat{\Delta}_n) \leq 0$, which implies a bound on $\|\hat{\Delta}_n\|_2$. Unlike previous analysis, the bound can be established without making any additional assumptions on the norm $R(\theta)$. We start with the following result, which expresses the upper bound on $\|\hat{\Delta}_n\|_2$ in terms of the gradient of the objective at θ^*.

Lemma 2 *Assume that the RSC condition is satisfied in E_r by the loss $\mathcal{L}(\cdot)$ with parameter κ. With $\hat{\Delta}_n = \hat{\theta}_{\lambda_n} - \theta^*$, for any norm $R(\cdot)$, we have*

$$\|\hat{\Delta}_n\|_2 \leq \frac{1}{\kappa}\|\nabla\mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*)\|_2 \,, \tag{22}$$

where $\nabla R(\cdot)$ is any sub-gradient of the norm $R(\cdot)$.

Figure 3 illustrates the above results. Note that the right hand side is simply the L_2 norm of the gradient of the objective evaluated at θ^*. For the special case when $\hat{\theta}_{\lambda_n} = \theta^*$, the gradient of the objective is zero, implying correctly that $\|\hat{\Delta}_n\|_2 = 0$. While the above result provides useful insights about the bound on $\|\hat{\Delta}_n\|_2$, the quantities on the right hand side depend on θ^*, which is unknown. We present another form of the result in terms of quantities such as λ_n, κ, and the norm compatibility constant $\psi(E_r) = \sup_{\mathbf{u}\in E_r} \frac{R(\mathbf{u})}{\|\mathbf{u}\|_2}$, which are often easier to compute or bound.

Theorem 2 *Assume that the RSC condition is satisfied in E_r by the loss $\mathcal{L}(\cdot)$ with parameter κ. With $\hat{\Delta}_n = \hat{\theta}_{\lambda_n} - \theta^*$, for any norm $R(\cdot)$, we have*

$$\|\hat{\Delta}_n\|_2 \leq \psi(E_r)\frac{1+\beta}{\beta}\frac{\lambda_n}{\kappa} \,. \tag{23}$$

The above result is deterministic, but contains λ_n and κ. In Section 3, we give precise characterizations of λ_n, which needs to satisfy (15). In Sections 4 and 6, we characterize the RSC condition constant κ for different losses and a variety of design matrices.

2.3 A Special Case: Decomposable Norms

In recent work, [29] considered regularized regression with the special case of decomposable norms, defined in terms of a pair of subspaces $\mathcal{M} \subseteq \bar{\mathcal{M}}$ of $\mathbb{R}^p$. The model is assumed to be in the subspace $\mathcal{M}$, and the

definition considers the so-called perturbation subspace $\bar{\mathcal{M}}^{\perp}$ which is the orthogonal complement of $\bar{\mathcal{M}}$. A norm $R(\cdot)$ is considered decomposable with respect to subspaces $(\mathcal{M}, \bar{\mathcal{M}}^{\perp})$ if $R(\theta+\gamma) = R(\theta) + R(\gamma)$ for all $\theta \in \mathcal{M}$ and $\gamma \in \bar{\mathcal{M}}^{\perp}$.

We show that for decomposable norms, the error set E_r in our analysis is included in the error cone defined in [29]. In the current context, let $\beta = 2$, $\theta^* \in \mathcal{M}$, then for any $\Delta = \Delta_{\bar{\mathcal{M}}^{\perp}} + \Delta_{\bar{\mathcal{M}}} \in E_r$, we have

$$
\begin{aligned}
& R(\theta^* + \Delta) && \leq\ R(\theta^*) + \frac{1}{2}R(\Delta) && (24)\\
\Rightarrow\quad & R(\theta^* + \Delta_{\bar{\mathcal{M}}^{\perp}} + \Delta_{\bar{\mathcal{M}}}) && \leq\ R(\theta^*) + \frac{1}{2}R(\Delta_{\bar{\mathcal{M}}^{\perp}} + \Delta_{\bar{\mathcal{M}}}) && (25)\\
\Rightarrow\quad & R(\theta^* + \Delta_{\bar{\mathcal{M}}^{\perp}}) - R(\Delta_{\bar{\mathcal{M}}}) && \overset{(a)}{\leq}\ R(\theta^*) + \frac{1}{2}R(\Delta_{\bar{\mathcal{M}}^{\perp}}) + \frac{1}{2}R(\Delta_{\bar{\mathcal{M}}}) && (26)\\
\Rightarrow\quad & R(\theta^*) + R(\Delta_{\bar{\mathcal{M}}^{\perp}}) - R(\Delta_{\bar{\mathcal{M}}}) && \overset{(b)}{\leq}\ R(\theta^*) + \frac{1}{2}R(\Delta_{\bar{\mathcal{M}}^{\perp}}) + \frac{1}{2}R(\Delta_{\bar{\mathcal{M}}}) && (27)\\
\Rightarrow\quad & R(\Delta_{\bar{\mathcal{M}}^{\perp}}) && \leq\ 3R(\Delta_{\bar{\mathcal{M}}}). && (28)
\end{aligned}
$$

where inequality (a) follows from the triangle inequality and (b) follows from decomposability of the norm. The last inequality is precisely the error cone in [29] for $\theta^* \in \mathcal{M}$. As a result, for any $\Delta \in E_r$, for decomposable norms we have

$$
R(\Delta) \quad = \quad R(\Delta_{\bar{\mathcal{M}}^{\perp}} + \Delta_{\bar{\mathcal{M}}}) \leq R(\Delta_{\bar{\mathcal{M}}^{\perp}}) + R(\Delta_{\bar{\mathcal{M}}}) \leq 4R(\Delta_{\bar{\mathcal{M}}}) \tag{29}
$$

Hence, the norm compatibility constant can be bounded as

$$
\psi(E_r) = \sup_{\Delta \in E_r} \frac{R(\Delta)}{\|\Delta\|_2} \leq 4 \sup_{\Delta \in E_r} \frac{R(\Delta_{\bar{\mathcal{M}}})}{\|\Delta\|_2} \leq 4 \sup_{\mathbf{u} \in \bar{\mathcal{M}} \setminus \{0\}} \frac{R(\mathbf{u})}{\|\mathbf{u}\|_2} = 4\Psi(\bar{\mathcal{M}}). \tag{30}
$$

where $\Psi(\bar{\mathcal{M}})$ is the subspace compatibility in $\bar{\mathcal{M}}$, as used in [29].

3 Bounds on the Regularization Parameter

Recall that the parameter λ_n needs to satisfy the inequality

$$
\lambda_n \geq \beta R^*(\nabla \mathcal{L}(\theta^*; Z^n)) \, . \tag{31}
$$

The right hand side of the inequality has two issues: the expression depends on θ^*, the optimal parameter which is unknown, and the expression is a random variable, since it depends on Z^n. In this section, we characterize the expectation $E[R^*(\nabla \mathcal{L}(\theta^*; Z^n))]$ in terms of the Gaussian width of the unit norm ball $\Omega_R = \{u : R(u) \leq 1\}$, and further discuss its upper bounds. For ease of exposition, we present results for the case of squared loss, i.e., $\mathcal{L}(\theta^*; Z^n) = \frac{1}{2n}\|y - X\theta^*\|^2$ with the linear model $y = X\theta + \omega$, where ω is noise vector with i.i.d. entries. Under this setting,

$$
\nabla \mathcal{L}(\theta^*; Z^n) = \frac{1}{n} X^T (y - X\theta^*) = \frac{1}{n} X^T \omega \, , \tag{32}
$$

which eliminates the dependency on the unknown θ^*. Before presenting the results, we introduce a few notations. We let $\Lambda_{\max}(\cdot)$ denote the largest eigenvalue of a square matrix. We also recall the definition of the sub-Gaussian norm for a sub-Gaussian variable x, $\|\|x\|\|_{\psi_2} = \sup_{p \geq 1} \frac{1}{\sqrt{p}} (E[|x|^p])^{1/p}$ [41].

From this section onwards, the analysis will take into account the randomness of the design X and the noise ω. Here we give a brief description of our assumptions on X and ω as follows,

Isotropic Sub-Gaussian Designs: the design matrix $X \in \mathbb{R}^{n\times p}$ has independent sub-Gaussian rows where each row satisfies $\|\|X_i\|\|_{\psi_2} \leq \kappa$ and $E[X_i X_i^T] = \mathbb{I}_{p\times p}$. Thus, the measure μ from which the rows X_i are sampled independently is an isotropic sub-Gaussian measure.
Anisotropic Sub-Gaussian Designs: the design matrix $X \in \mathbb{R}^{n\times p}$ has independent rows, and each row X_i is anisotropic sub-Gaussian with $E[X_i^T X_i] = \Sigma$. Further, we assume that corresponding isotropic random vector $\tilde{X}_i = X_i\Sigma^{-1/2}$ satisfies $\left\|\left\|\tilde{X}_i\right\|\right\|_{\psi_2} \leq \kappa$. A simple special case of such an anisotropic sub-Gaussian design is when $X_i \sim N(0,\Sigma)$, where $\tilde{X}_i = X_i\Sigma^{-1/2} \sim N(0,\mathbb{I})$ so that $\left\|\left\|\tilde{X}_i\right\|\right\|_{\psi_2} = 1$.
Sub-Gaussian Noise: the noise ω has i.i.d. centered unit-variance sub-Gaussian entries with $\|\|\omega_i\|\|_{\psi_2} \leq K$.
For convenience, we only use the shorthand in bold font to specify the assumptions. In the following theorem, we characterize the expectation of $R^*(\nabla\mathcal{L}(\theta^*; Z^n))$ in terms of Gaussian width of the unit norm ball $w(\Omega_R)$.

Theorem 3 *Let $\Omega_R = \{u : R(u) \leq 1\}$, and $\mathcal{L}$ be the squared loss. For sub-Gaussian design X and noise ω, we have*

$$E\left[R^*(\nabla\mathcal{L}(\theta^*; Z^n))\right] \leq \eta\xi \cdot \frac{\kappa w(\Omega_R)}{\sqrt{n}} , \tag{33}$$

where the expectation is taken over both X and ω. The constant ξ is given by

$$\xi = \begin{cases} 1 & \text{if } X \text{ is isotropic} \\ \sqrt{\Lambda_{\max}(\Sigma)} & \text{if } X \text{ is anisotropic} . \end{cases}$$

Bounding the expectation of $R^*(\nabla\mathcal{L}(\theta^*; Z^n))$ gives us a rough scale of the regularization parameter λ_n. In the next theorem, we present a high-probability upper bound for $R^*(\nabla\mathcal{L}(\theta^*; Z^n))$.

Theorem 4 *Let design X and noise ω be sub-Gaussian, and $\mathcal{L}$ be squared loss. Define $\phi = \sup_{R(u)\leq 1}\|u\|_2$, then for any $\tau > 0$, with probability at least $1 - c_1\exp\left(-\min\left((\frac{\tau}{c_2\xi\kappa\phi})^2, c_0 n\right)\right)$, we have*

$$R^*\left(\nabla\mathcal{L}(\theta^*; Z^n)\right) \leq \sqrt{\frac{2K^2+1}{n}}\left(c\xi\kappa \cdot w(\Omega_R) + \tau\right) , \tag{34}$$

where c, c_0, c_1 and c_2 are all absolute constants, and ξ is the same as in Theorem 3.

Bounding the Gaussian width $w(\Omega_R)$: In certain cases, one may be able to directly obtain a bound on the Gaussian width $w(\Omega_R)$. Here, we provide a mechanism for bounding the Gaussian width $w(\Omega_R)$ of the unit norm ball in terms of the Gaussian width of a suitable cone, obtained by shifting or translating the norm ball. In particular, the result involves taking any point on the boundary of the unit norm ball, considering that as the origin, and constructing a cone using the norm ball. Since such a construction can be done with any point on the boundary, the tightest bound is obtained by taking the infimum over all points on the boundary. The motivation behind getting an upper bound of the Gaussian width $w(\Omega_R)$ of the unit norm ball in terms of the Gaussian width of such a cone is because considerable advances have been made in recent years in upper bounding Gaussian widths of such cones [14, 1].

Lemma 3 *Let $\Omega_R = \{u : R(u) \leq 1\}$ be the unit norm ball and $\Theta_R = \{u : R(u) = 1\}$ be the boundary. For any $\tilde{\theta} \in \Theta_R$, $\rho(\tilde{\theta}) = \sup_{\theta: R(\theta)\leq 1}\|\theta - \tilde{\theta}\|_2$ is the diameter of Ω_R measured with respect to $\tilde{\theta}$. Let $G(\tilde{\theta}) = \mathrm{cone}(\Omega_R - \tilde{\theta}) \cap \rho(\tilde{\theta})B_2^p$, i.e., the cone of $(\Omega_R - \tilde{\theta})$ intersecting the ball of radius $\rho(\tilde{\theta})$. Then*

$$w(\Omega_R) \leq \inf_{\tilde{\theta}\in\Theta_R} w(G(\tilde{\theta})) . \tag{35}$$

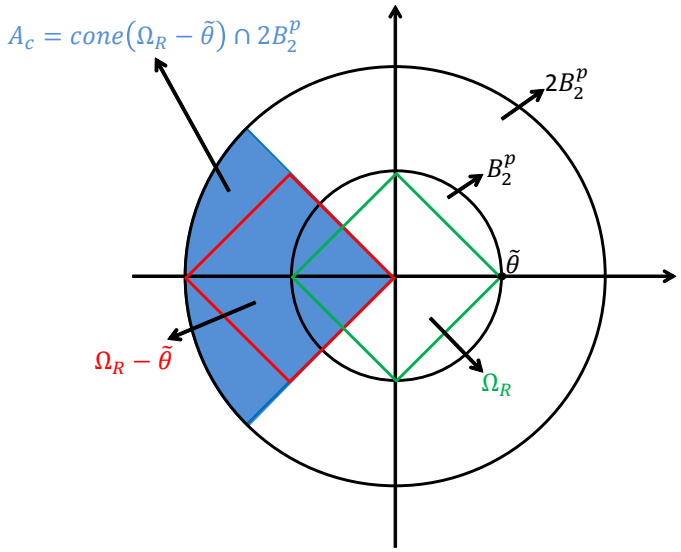

Figure 4: Bounding the Gaussian width of a norm ball, e.g., corresponding to L_1 norm, by shifting the norm ball and using the width of the corresponding cone (Lemma 3). The approach allows one to directly use existing results on bounding Gaussian widths of certain cones. In some cases, it may be easier to directly bound the Gaussian width of the norm ball, rather than using the shifting argument.

The analysis and results for λ_n presented above can be extended to general convex losses arising in the context of GLMs for sub-Gussian designs and sub-Gaussian noise (see Section 6). .

4 Least Squares Models: Restricted Eigenvalue Conditions

The error bound analysis in Theorem 2 depends on the restricted strong convexity (RSC) assumption. In this section, we establish RSC conditions for sub-Gaussian design matrices when the loss function is the squared loss. For squared loss, i.e., $\mathcal{L}(\theta; Z^n) = \frac{1}{n}\|\mathbf{y} - X\theta\|^2$, the RSC condition (20) becomes equivalent to the Restricted Eigenvalue (RE) condition [6, 29], since

$$\begin{aligned}\delta\mathcal{L}(\Delta, \theta^*) &= \frac{1}{n}\|y - X(\theta^* + \Delta)\|^2 - \frac{1}{n}\|y - X\theta^*\|^2 + \frac{1}{n}\langle X^T(y - X\theta^*), \Delta\rangle \\ &= \frac{1}{n}\|X\Delta\|^2 = \frac{1}{n}\sum_{i=1}^{n}\langle X_i, \Delta\rangle^2 \,.\end{aligned} \tag{36}$$

so that the condition simplifies to

$$\delta\mathcal{L}(\Delta, \theta^*) = \frac{1}{n}\sum_{i=1}^{n}\langle X_i, \Delta\rangle^2 \geq \kappa\|\Delta\|_2^2 \,, \tag{37}$$

for all $\Delta \in E_r$. We make two simplifications which lets us develop the RE results in terms of widths of spherical caps rather than over the error set E_r. Let n_{E_r} be the sample complexity for the RE condition over the set E_r, so that for $n > n_{E_r}$ samples, with high probability

$$\inf_{\Delta \in E_r} \frac{1}{n}\|X\Delta\|_2^2 \geq \kappa_{E_r}\|\Delta\|_2^2 \,, \tag{38}$$

for some $\kappa_{E_r} > 0$. Let $C_r = \text{cone}(E_r)$ and let n_{C_r} be the sample complexity for the RE condition over the cone C_r, so that for $n > n_{C_r}$ samples, with high probability

$$\inf_{\Delta \in C_r} \frac{1}{n}\|X\Delta\|_2^2 \geq \kappa_{C_r}\|\Delta\|_2^2 \,, \tag{39}$$

for some $\kappa_{C_r} > 0$. Since $E_r \subseteq C_r$, we have $n_{E_r} \leq n_{C_r}$. Thus, it is sufficient to obtain (an upper bound to) the sample complexity n_{C_r}, since that will serve as an upper bound to n_{E_r}, the sample complexity over E_r. Further, since C_r is a cone, the absolute magnitude $\|\Delta\|_2$ does not affect the sample complexity. As a result, it is sufficient to focus on a spherical cap $A = C_r \cap S^{p-1}$. In particular, if n_A denotes the sample complexity for the RE condition over the spherical cap A, so that for $n > n_A$ samples, with high probability

$$\inf_{u \in A} \frac{1}{n}\|Xu\|_2^2 \geq \bar{\kappa}_A\|u\|_2^2 \,, \tag{40}$$

for some $\bar{\kappa}_A > 0$, then $n_A = n_{C_r} \geq n_{E_r}$. Noting that $\|u\|_2 = 1$ for $u \in C_r \cap S^{p-1}$, we consider sample complexity results for RE conditions the form

$$\inf_{u \in A} \frac{1}{n}\|Xu\|_2^2 = \inf_{u \in A} \frac{1}{n}\sum_{i=1}^{n}\langle X_i, u\rangle^2 \geq \kappa_A(n,p) \tag{41}$$

where $\kappa_A(n,p) > 0$ with high probability for $n > n_A$. In this section, we characterize sample complexity n_A over *any* given spherical cap A, and establish RE conditions for isotropic and anisotropic sub-Gaussian design matrices X in terms of the Gaussian width $w(A)$.

Analysis of RE conditions for certain types of design matrices for certain types of norms, especially the L_1 norm, have appeared in the literature [6, 33, 34]. The RE/RIP conditions for independent isotropic Gaussian designs have been widely studied for the case of L_1 norm [10, 6]. The generalization to RE condition for correlated Gaussian designs for the special of L_1 norm was studied in [33]. [14] consider the more general context of atomic norms, and RE condition analysis applies to any spherical cap A, with sample complexity results in terms of $w(A)$, the Gaussian width of A. However, the analysis relies on Gordon's inequality [19, 20, 24], which is applicable only for isotropic Gaussian design matrices. Progress has been made on establishing RE conditions for sub-Gaussian designs for error sets/caps corresponding to specific norms such as L_1 [46]. In recent work, RE conditions were developed for anisotropic sub-Gaussian designs for the L_1 norm [34]. Further, recent work have pointed out the differences between the RE and the RIP condition, which gives a two-sided bound on quadratic forms of random matrices [29]. In particular, while the RE condition is sufficient for structured estimation, the RIP results are stronger and may have higher sample complexity.

In the following, we establish the stronger RIP results for any spherical cap A and any sub-Gaussian design matrix, handling the isotropic and anisotropic cases separately. The special case of Gaussian design matrices are automatically covered by the sub-Gaussian results, and results such as Gordon's inequality can be viewed as a special case. All results are in terms of $w(A)$, the Gaussian width of A, even for sub-Gaussian designs. In fact, all existing RE results do implicitly have the width term, but in a form specific to the chosen norm [33, 34]. The analysis on atomic norm in [14] has the $w(A)$ term explicitly, but the analysis relies on Gordon's inequality [19, 20, 24], which is applicable only for isotropic Gaussian design matrices.

The proof technique we use is an application of generic chaining [37, 38]. The specific form we utilize was originally developed in [21, 28]. The main idea is to pose the RIP condition as a bound on the supremum of a suitable stochastic process, so that generic chaining can be invoked to obtain a bound. The key difference between our RIP analysis and much of the existing literature on RE conditions, which use specialized tools such as Gaussian comparison principles [33, 29] or analysis geared to a particular norm [34], is the use

of generic chaining which simplifies the analysis considerably. Further, the RIP results can be viewed as a generalization of the celebrated Johnson-Lindenstrauss (JL) lemma [17], and the interested reader can explore these connections in [21].

Isotropic Sub-Gaussian Designs: We first consider the case where the design matrix $X \in \mathbb{R}^{n\times p}$ has independent sub-Gaussian rows where each row satisfies $\|\|X_i\|\|_{\psi_2} \leq \kappa$ and $E[X_i X_i^T] = \mathbb{I}_{p\times p}$. Thus, the measure μ from which the rows X_i are sampled independently is an isotropic sub-Gaussian measure.

Theorem 5 *Let X be a design matrix with independent isotropic sub-Gaussian rows, i.e., $\|\|X_i\|\|_{\psi_2} \leq \kappa$ and $E[X_i X_i^T] = \mathbb{I}_{p\times p}$. Then, for absolute constants $\eta, c > 0$, with probability at least $(1 - 2\exp(-\eta w^2(A)))$, we have*

$$\sup_{u\in A}\left|\frac{1}{n}||Xu||^2 - 1\right| \;=\; \sup_{u\in A}\left|\frac{1}{n}\sum_{i=1}^{n}\langle X_i, u\rangle^2 - 1\right| \;\leq\; c\frac{w(A)}{\sqrt{n}}\,, \tag{42}$$

or, equivalently,

$$1 - c\frac{w(A)}{\sqrt{n}} \;\leq\; \inf_{u\in A}\frac{1}{n}||Xu||^2 \;\leq\; \sup_{u\in A}\frac{1}{n}||Xu||^2 \;\leq\; 1 + c\frac{w(A)}{\sqrt{n}}\,. \tag{43}$$

As a result, for $n > c^2 w^2(A)$, the RE condition: $\inf_{u\in A} \|Xu\|^2 \geq 1 - cw(A)/\sqrt{n} > 0$ is satisfied with high probability for any sub-Gaussian design matrix. More generally, choosing $\epsilon = cw(A)/\sqrt{n}$, one can write the result in a traditional RIP form [10].

Anisotropic Sub-Gaussian Designs: We now consider the case where the design matrix $X \in \mathbb{R}^{n\times p}$ has independent rows, and each row X_i is anisotropic sub-Gaussian with $E[X_i^T X_i] = \Sigma$. Further, we assume that corresponding isotropic random vector $\tilde{X}_i = X_i\Sigma^{-1/2}$ satisfies $\left\|\left\|\tilde{X}_i\right\|\right\|_{\psi_2} \leq \kappa$. A simple special case of such an anisotropic sub-Gaussian design is when $X_i \sim N(0, \Sigma)$, where $\tilde{X}_i = X_i\Sigma^{-1/2} \sim N(0, \mathbb{I})$ so that $\left\|\left\|\tilde{X}_i\right\|\right\|_{\psi_2} = 1$. The result below characterizes RIP-style property of any such anisotropic sub-Gaussian designs.

Theorem 6 *Let X be a design matrix with independent anisotropic sub-Gaussian rows, i.e., $E[X_i^T X_i] = \Sigma$ and $\left\|\left\|X_i\Sigma^{-1/2}\right\|\right\|_{\psi_2} \leq \kappa$. Then, for absolute constants $\eta, c > 0$, with probability at least $(1-2\exp(-\eta w^2(A)))$, we have*

$$\sup_{u\in A}\left|\frac{1}{n}\frac{1}{u^T\Sigma u}||Xu||^2 - 1\right| \;=\; \sup_{u\in A}\left|\frac{1}{n}\frac{1}{u^T\Sigma u}\sum_{i=1}^{n}\langle X_i, u\rangle^2 - 1\right| \;\leq\; c\frac{w(A)}{\sqrt{n}}\,. \tag{44}$$

Further,

$$\lambda_{\min}(\Sigma|A)\left(1 - c\frac{w(A)}{\sqrt{n}}\right) \;\leq\; \inf_{u\in A}\frac{1}{n}||Xu||^2 \;\leq\; \sup_{u\in A}\frac{1}{n}||Xu||^2 \;\leq\; \lambda_{\max}(\Sigma|A)\left(1 + c\frac{w(A)}{\sqrt{n}}\right)\,, \tag{45}$$

where

$$\lambda_{\min}(\Sigma|A) = \inf_{u\in A}\; u^T\Sigma u\,, \qquad \textit{and} \qquad \lambda_{\max}(\Sigma|A) = \sup_{u\in A}\; u^T\Sigma u \tag{46}$$

are the restricted minimum and maximum eigenvalues of Σ restricted to $A \subseteq S^{p-1}$.

Thus, for the anisotropic case, the RIP is with respect to the restricted minimum and maximum eigenvalues corresponding to $A \subseteq S^{p-1}$. For the special case when $A = S^{p-1}$, we have $\lambda_{\min}(\Sigma|A) = \lambda_{\min}(\Sigma)$, the minimum eigenvalue, and $\lambda_{\max}(\Sigma|A) = \lambda_{\max}(\Sigma)$, the maximum eigenvalue of Σ. Further, when $\Sigma = \mathbb{I}$,

we get back the result in Theorem 11. Finally, it is instructive to compare the above result to existing characterizations of the RE condition for anisotropic Gaussian [33] and anisotropic sub-Gaussian [34] designs, focused on the L_1 norm. The result in Theorem 12 is a more general RIP result, applies to any spherical cap A, and is in terms of $w(A)$, the Gaussian width of the spherical cap A.

5 Examples and Applications

In this section, we give examples of the analysis from previous sections for three norms: L_1 norm, group sparse norm, and L_2 norm. The summary of the results is given in Table 1. Other examples can be constructed for norms and error sets with known bounds on the Gaussian widths and norm compatibility constants [14, 29], and more general ways to bound the Gaussian widths and norm compatibility constants have been developed in [25, 16].

$\mathbf{L_1}$ Norm: Assume that the statistical parameter θ^* is s-sparse, and note that $\|\theta^*\|_1 \leq \sqrt{s}\|\theta^*\|_2$. Since L_1 norm is a decomposable norm following the result in (28), we have $\Psi(E_r) \leq 4\Psi(\bar{\mathcal{M}}) = 4\sqrt{s}$.

Applying Lemma 3, let $\bar{\theta}$ be a 1-sparse vector, and $\rho(\bar{\theta}) = 2$, then $w(\Omega_R)$ can be bounded by

$$w(\Omega_R) \leq \inf_{\tilde{\theta} \in \Theta_R} w(G(\tilde{\theta})) = w(G(\bar{\theta})) \overset{(a)}{=} O\left(\sqrt{\log p}\right), \tag{47}$$

where (a) is obtained from the fact that Gaussian width of $G(\tilde{\theta})$ with $\tilde{\theta}$ be a s-sparse vector is $\sqrt{2s\log(\frac{p}{s}) + \frac{5}{4}s}$ [14]. See Figure 4 for more details. From Theorem 4 and (47), the bound on λ_n is

$$\lambda_n \leq c\frac{w(\Omega_R)}{\sqrt{n}} = O\left(\sqrt{\frac{\log p}{n}}\right). \tag{48}$$

Hence, the recovery error is bounded by

$$\|\hat{\Delta}_n\|_2 \leq c_3\frac{\Psi(E_r)\lambda_n}{\kappa} = O\left(\sqrt{\frac{s\log p}{n}}\right), \tag{49}$$

which is similar to the results obtained in well known results [14, 29].

Group Sparse Norm: Suppose that the index set $\{1, 2, \cdots, p\}$ can be partitioned into a set of T disjoint groups, say $\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2, \cdots, \mathcal{G}_T\}$. Define $(1, \nu)$-group norm for a given vector $\boldsymbol{\nu} = (\nu_1, \cdots, \nu_T) \in [1, \infty]^T$ as

$$\|\alpha\|_{\mathcal{G},\boldsymbol{\nu}} = \sum_{t=1}^{T} \|\alpha_{\mathcal{G}_t}\|_{\nu_t} \tag{50}$$

As shown in [29] Group norm is a decomposable norm. For a given subset $S_{\mathcal{G}} \subset \{1, \ldots, T\}$ with cardinality $|S_{\mathcal{G}}|$, define the subspace $A(S_{\mathcal{G}}) = \{\alpha \in \mathbb{R}^p \,|\alpha_{\mathcal{G}_t} = 0, \;\; \forall t \notin S_{\mathcal{G}}\}$. Let $\nu_t \geq 2$, then we have

$$\|\Delta\|_{\mathcal{G},\boldsymbol{\nu}} = \sum_{t \in S_{\mathcal{G}}} \|\Delta_{\mathcal{G}_t}\|_{\nu_t} \leq \sum_{t \in S_{\mathcal{G}}} \|\Delta_{\mathcal{G}_t}\|_2 \leq \sqrt{s_{\mathcal{G}}}\|\Delta\|_2. \tag{51}$$

Hence, from (30) and (51) we have

$$\Psi(E_r) \leq 4\sqrt{s_{\mathcal{G}}}. \tag{52}$$

Applying Lemma 3, define $\bar{\theta}$ with 1-active group, and $\rho(\bar{\theta}) = 2$, then $w(\Omega_R)$ can be bounded by

$$w(\Omega_R) \leq \inf_{\tilde{\theta} \in \Theta_R} w(G(\tilde{\theta})) = w(G(\bar{\theta})) \overset{(a)}{=} O\left(\sqrt{m + \log T}\right), \tag{53}$$

$R(u)$	$\lambda_n := c_1 \frac{w(\Omega_R)}{\sqrt{n}}$	$\kappa := \left[\max\left\{\left(1-\sqrt{c_2}\frac{w(A)}{\sqrt{n}}\right), 0\right\}\right]^2$	$\Psi(E_r)$	$\|\hat{\Delta}_n\|_2 := c_3 \frac{\Psi(E_r)\lambda_n}{\kappa}$
L_1	$O\left(\sqrt{\frac{\log p}{n}}\right)$	$\Theta(1)$ if $n > c_2 w^2(A) = O(s \log p)$	$\sqrt{s}$	$O\left(\sqrt{\frac{s \log p}{n}}\right)$
Group sparse	$O\left(\sqrt{\frac{m+\log T}{n}}\right)$	$\Theta(1)$ if $n > c_2 w^2(A) = O(s_{\mathcal{G}}(m + \log T))$	$\sqrt{s_{\mathcal{G}}}$	$O\left(\sqrt{\frac{s_{\mathcal{G}}(m+\log T)}{n}}\right)$
L_2	$O\left(\sqrt{\frac{p}{n}}\right)$	$\Theta(1)$ if $n > c_2 w^2(A) = O(p)$	1	$O\left(\sqrt{\frac{p}{n}}\right)$

Table 1: A summary of values for the regularization parameter λ_n, the RE condition constant κ, the norm constant $\Psi(E_r)$ and recovery bounds $\|\hat{\Delta}_n\|_2$ for ℓ_1, ℓ_2 and group norms in case of Gaussian Design matrix with Gaussian noise. All results are given up to constants with more emphasis on the scale of the results.

where $m = \max_t |\mathcal{G}_t|$ and (a) is obtained from the fact that Gaussian width of $G(\tilde{\theta})$ where $\tilde{\theta}$ has k active group is $\sqrt{2k(m + \log(T-k)) + k}$ [14]. From Theorem 4 and (53), the bound on λ_n is

$$\lambda_n \leq c\frac{w(\Omega_R)}{\sqrt{n}} = O\left(\sqrt{\frac{m+\log T}{n}}\right). \tag{54}$$

Hence, the recovery error is bounded by

$$\|\hat{\Delta}_n\|_2 \leq c_3 \frac{\Psi(E_r)\lambda_n}{\kappa} = O\left(\sqrt{\frac{s_{\mathcal{G}}(m+\log T)}{n}}\right), \tag{55}$$

which is similar to the results obtained in previous works [14, 29].

$\mathbf{L_2}$ Norm: With L_2 norm as the regularizer, the norm constant is obtained as

$$\Psi(E_r) = \sup_{\Delta \in E_r} \frac{\|\Delta\|_2}{\|\Delta\|_2} = 1. \tag{56}$$

Applying Lemma 3, set $\rho(\tilde{\theta}) = 1$, then $w(\Omega_R)$ can be bounded by

$$w(\Omega_R) \leq \inf_{\tilde{\theta} \in \Theta_R} w(G(\tilde{\theta})) = O\left(\sqrt{p}\right). \tag{57}$$

From Theorem 4 and (57), the bound on λ_n is

$$\lambda_n \leq c\frac{w(\Omega_R)}{\sqrt{n}} = O\left(\sqrt{\frac{p}{n}}\right). \tag{58}$$

Hence, the recovery error is bounded by

$$\|\hat{\Delta}_n\|_2 \leq c_3 \frac{\Psi(E_r)\lambda_n}{\kappa} = O\left(\sqrt{\frac{p}{n}}\right). \tag{59}$$

6 Generalized Linear Models: Restricted Strong Convexity

In this section, we extend our results to estimation with norm regularization in the context of generalized linear models (GLMs) [4, 8]. Assume that the conditional distribution of the response y_i conditioned on the covariates X_i is an exponential family distribution:

$$p(y_i|X_i;\theta^*) = p(y_i|\langle X_i,\theta^*\rangle) = \exp\{y_i\langle X_i,\theta^*\rangle - \varphi(\langle X_i,\theta^*\rangle)\}\ . \tag{60}$$

where

$$\varphi(\langle X_i,\theta^*\rangle) = \log\left(\int_{y_i}\exp\{y_i\langle X_i,\theta^*\rangle\}\ dy_i\right) \tag{61}$$

is the log-partition function [8, 4, 44].[2] In GLMs, the conditional distribution of the response y_i is characterized by an exponential family distribution $p(y_i|\eta_i)$ with natural parameter $\eta_i = \langle X_i,\theta^*\rangle$ determined by the covariates X_i and the parameter θ^*. It is easy to verify that the gradient of the log-partition function w.r.t. the natural parameter $\eta_i = \langle X_i,\theta^*\rangle$ gives the expectation of the response [4, 8], i.e.,

$$\nabla_{\eta_i}\varphi(\eta_i) = \nabla_{\langle X_i,\theta^*\rangle}\varphi(\langle X_i,\theta^*\rangle) = E[y_i|\langle X_i,\theta^*\rangle]\ . \tag{62}$$

For estimating θ^*, the loss function corresponding to GLMs typically consider the negative log likelihood of the conditional distribution:

$$\mathcal{L}(\theta;Z^n) = -\frac{1}{n}\log p(y_i|X_i;\theta^*) = \frac{1}{n}\sum_{i=1}^{n}(\varphi(\langle X_i,\theta\rangle) - y_i\langle X_i,\theta\rangle)\ . \tag{63}$$

In the current context, we assume $\theta*$ to be sparse/structured, and the structure can be suitably captured by a norm $R(\cdot)$. Then, the estimation of θ^* with norm regularization takes the form:

$$\hat{\theta}_{\lambda_n} = \operatorname*{argmin}_{\theta\in\mathbb{R}^p}\mathcal{L}(\theta;Z^n) + \lambda_n R(\theta) = \operatorname*{argmin}_{\theta\in\mathbb{R}^p}\frac{1}{n}\sum_{i=1}^{n}(\varphi(\langle X_i,\theta\rangle) - y_i\langle X_i,\theta\rangle) + \lambda_n R(\theta)\ . \tag{64}$$

Noise in the context of GLMs is simply the deviation of a specific response y_i from the conditional mean, i.e., $\omega_i = E[y_i|X_i] - y_i$. Popular examples of GLMs come from suitable choices of the conditional distribution, e.g., when $p(y_i|\langle X_i,\theta\rangle)$ is Gaussian so that $\varphi(\langle X_i,\theta\rangle) = \frac{\langle X_i,\theta\rangle^2}{2}$, Bernoulli so that $\varphi(\langle X_i,\theta\rangle) = \log(1+\exp(\langle X_i,\theta\rangle))$, and Poisson where $\varphi(\langle X_i,\theta\rangle) = \exp(\langle X_i,\theta\rangle)$, respectively yielding least squares regression, logistic regression, and Poisson regression loss functions. Next, we provide the key results needed to characterize the regularization parameter λ_n and restricted strong convexity (RSC) in the context of GLMs. The non-asymptotic bound on the estimation error then follows from the general result in (23).

Bounds on the Regularization Parameter: Following the general analysis from Section 3, the regularization parameter needs to satisfy the condition: $\lambda_n \geq \beta R^*(\nabla_\theta\mathcal{L}(\theta^*;Z^n))$ for any fixed $\beta > 1$. For GLMs,

$$\nabla_\theta\mathcal{L}(\theta^*;Z^n) = -\frac{1}{n}\sum_{i=1}^{n}y_iX_i + \frac{1}{n}\sum_{i=1}^{n}X_i\nabla_{\langle X_i,\theta^*\rangle}\varphi(\langle X_i,\theta^*\rangle) = \frac{1}{n}\sum_{i=1}^{n}X_i(E[y|X_i] - y_i) = \frac{1}{n}X^T\omega\ , \tag{65}$$

where we have used the fact, $\nabla_{\langle X_i,\theta^*\rangle}\varphi(\langle X_i,\theta^*\rangle) = E[y_i|\langle X_i,\theta^*\rangle]$ and $\omega = E[y_i|\langle X_i,\theta^*\rangle] - y_i$. Thus, the form of $\nabla_\theta\mathcal{L}(\theta^*;Z^n)$ is the same as that in Section 3. Assuming the design matrix X and noise ω are sub-Gaussian, a characterization of λ_n follows from Theorems 3 and 4 in Section 3. In particular, $E[\nabla_\theta\mathcal{L}(\theta^*;Z^n)] = O(\frac{w(\Omega_R)}{\sqrt{n}})$, with corresponding high probability concentration results, and it suffices to have λ_n to be of this order.

[2] Note that for GLMs over discrete responses y_i, the integration needs to be suitably changed to summation [4, 8].

Restricted Strong Convexity: By definition, the restricted strong convexity considers

$$\delta\mathcal{L}(u,\theta^*) = \mathcal{L}(\theta^*+u) - \mathcal{L}(\theta^*) - \langle \nabla\mathcal{L}(\theta^*), u\rangle = \frac{1}{n}\sum_{i=1}^{n} \nabla^2\varphi(\langle\theta^*, X_i\rangle + \gamma_i\langle u, X_i\rangle)\langle u, X_i\rangle^2 \,,$$

where $\gamma_i \in [0,1]$, and where the last equality follows from a direct application of the mean value theorem [35]. Since the log-partition function φ is of Legendre type [3, 4, 8], the second derivative $\nabla^2\varphi(\cdot)$ is always positive. Since the RSC condition relies on a non-trivial lower bound for the above quantity, the analysis will consider suitable compact sets where $\nabla^2\varphi(\cdot)$ is bounded away from zero by a constant. In particular, for a suitable constant T, we consider the sets $\{X_i | |\langle X_i, \theta^*\rangle| < T\}$ and $\{X_i | |\langle X_i, u\rangle| < T\}$. For X_i lying in these sets, the argument $a = \langle X_i, \theta^*\rangle + \gamma_i\langle u, X_i\rangle$ of the second derivative satisfies $|a| \le 2T$, which is the compact set of interest. Within the compact set, $\ell = \ell_\varphi(T) = \min_{|a|\le 2T} \nabla^2\varphi(a)$ is bounded away from zero. Outside the compact set, we will only assume $\nabla^2\varphi(\cdot) > 0$. Based on the above construction, we have

$$\delta\mathcal{L}(u,\theta^*) \ge \frac{\ell}{n}\sum_{i=1}^{n} \langle X_i, u\rangle^2 \, \mathbb{I}[|\langle X_i, \theta^*\rangle| < T] \, \mathbb{I}[|\langle X_i, u\rangle| < T] \,. \tag{66}$$

The quadratic form based lower bound allows us to establish RSC conditions for GLMs with isotropic subGaussian design matrices by building on results from Section 4 for RE conditions for squared loss. As a result, the sample complexity of the RSC condition is also expressed in terms of the Gaussian width of the spherical cap A derived from the error set. The analysis can be suitably generalized to anisotropic design matrices using techniques discussed in Section 4.

As before, we consider $u \in A \subseteq S^{p-1}$ so that $\|u\|_2 = 1$. Assuming X has isotropic sub-Gaussian rows with $\|\!|X_i|\!\|_{\varphi_2} \le \kappa$, $\langle X_i, \theta^*\rangle$ and $\langle X_i, u\rangle$ are sub-Gaussian random variables with sub-Gaussian norm at most $C\kappa$ [42]. Denote by ε_1 and ε_2 the probability that $\langle X_i, u\rangle$ and $\langle X_i, \theta^*\rangle$ exceeds some constant T, i.e., $\varepsilon_1(T;u) = P\{|\langle X_i, u\rangle| > T\} \le e\cdot\exp(-c_2T^2/C^2\kappa^2) = \bar{\varepsilon}_1$, and $\varepsilon_2(T;\theta^*) = P\{|\langle X_i, \theta^*\rangle| > T\} \le e\cdot\exp(-c_2T^2/C^2\kappa^2) = \bar{\varepsilon}_2$, where $\bar{\varepsilon}_1 = \bar{\varepsilon}_1(T,\kappa)$ and $\bar{\varepsilon}_2 = \bar{\varepsilon}_2(T;\kappa)$ are uniform upper bounds on the individual tail probabilities. The result we present below is in terms of the above defined constants $\ell = \ell_\varphi(T)$, $\bar{\varepsilon}_1 = \bar{\varepsilon}_1(T,\kappa)$ and $\bar{\varepsilon}_2 = \bar{\varepsilon}_2(T,\kappa)$ for any suitably chosen T.

Theorem 7 *Let $X \in \mathbb{R}^{n\times p}$ be a design matrix with independent isotropic sub-Gaussian rows such that $\|\!|X_i|\!\|_{\varphi_2} \le \kappa$. Then, for any set $A \subseteq S^{p-1}$ for suitable constants $\eta, c > 0$, with probability at least $1 - 2\exp\big(-\eta w^2(A)\big)$, we have*

$$\inf_{u\in A} \partial\mathcal{L}(u,\theta^*) \ge \ell\underline{\rho}^2\left(1 - c\kappa_1^2\frac{w(A)}{\sqrt{n}}\right) \,, \tag{67}$$

where $\underline{\rho}^2 = \inf_{u\in A}\rho_u^2$, with $\rho_u^2 = E[\langle X_i, u\rangle^2\mathbb{I}[|\langle X_i, \theta^*\rangle| < T]\mathbb{I}[|\langle X_i, u\rangle| < T]]$, and $\kappa_1 = \frac{\kappa}{1-\bar{\varepsilon}_1-\bar{\varepsilon}_2}$.

The form of the result is closely related to the corresponding result for the RE condition $\inf_{u\in A}\|Xu\|_2$ considered in Section 4. Note that RSC analysis for GLMs was considered in [29] for specific norms, especially L_1, whereas our analysis applies to any set $A \subseteq S^{p-1}$, hence to any norm, and the result is in terms of the Gaussian width $w(A)$ of A. Further, following arguments in Section 4, the RE analysis for GLMs can be extended to anisotropic subGaussian design matrices.

7 Conclusions

The paper presents a general set of results and tools for characterizing non-asymptotic estimation error in norm regularized regression problems. The analysis holds for any norm, and subsumes much of existing

literature focused on structured sparsity and related themes. The work can be viewed as a direct generalization of results in [29], which presented related results for decomposable norms. Our analysis illustrates the important role Gaussian widths, as a measure of size of suitable sets, play in such results. Further, the error sets for regularized and constrained versions of such problems are shown to be closely related [6].

While the paper presents a unified geometric treatment of non-asymptotic structured estimation with regularized estimators, several technical questions need further investigation. The focus of the analysis has been on thin-tailed distributions, and the RE/RSC type analysis presented really gives two sided bounds, i.e., RIP, showing that thin-tailed distributions do satisfy the RIP condition. For heavy tailed measurements, the lower and upper tails of quadratic forms behave differently [30, 27], and it may be possible to establish geometric estimation error analysis for general norms, some special cases of which have been investigated in recent years [22, 23, 27]. Further, the sample complexity of the phase transitions in the RE/RSC conditions for anisotropic designs depend on the largest eigenvalue (operator norm) of the covariance matrix, making the estimator sample inefficient for highly correlated designs. Since real-world several problems, including spatial and temporal problems, do have correlated observations, it will be important to investigate estimators which perform well in such settings [18]. Finally, the focus of the work is on parametric estimation, and it will be interesting to explore generalizations of the analysis to non-parametric settings.

Appendix

A Background and Preliminaries

We start with a review of some definitions and well-known results which will be used for our proofs.

A.1 Gaussian Width

In several of our proofs, we use the concept of Gaussian width [20, 14], which is defined as follows.

Definition 1 (Gaussian width) For any set $A \in \mathbb{R}^p$, the *Gaussian width* of the set A is defined as:

$$w(A) = E_g \left[\sup_{u \in A} \langle g, u \rangle \right] , \tag{68}$$

where the expectation is over $g \sim N(0, \mathbb{I}_{p \times p})$, a vector of independent zero-mean unit-variance Gaussian random variable.

The Gaussian width $w(A)$ provides a geometric characterization of the size of the set A. We consider three perspectives of the Gaussian width, and provide some properties which are used in our analysis. First, consider the Gaussian process $\{Z_u\}$ where the constituent Gaussian random variables $Z_u = \langle t, g \rangle$ are indexed by $u \in A$, and $g \sim N(0, \mathbb{I}_{p \times p})$. Then the Gaussian width $w(A)$ can be viewed as the expectation of the supremum of the Gaussian process $\{Z_t\}$. Bounds on the expectations of Gaussian and other empirical processes have been widely studied in the literature, and we will make use of generic chaining for some of our analysis [37, 38, 7, 24]. Second, $\langle u, g \rangle$ can be viewed as a Gaussian random projection of each $u \in A$ to one dimension, and the Gaussian width simply measures the expectation of largest value of such projections. Third, if A is the unit ball of any norm $R(\cdot)$, i.e., $A = \{x \in \mathbb{R}^p \mid R(x) \leq 1\}$, then $w(A) = E_g[R^*(g)]$ by definition of the dual norm. Thus, the Gaussian width is the expected value of the dual norm of a standard Gaussian random vector. For instance, if A is unit ball of L_1 norm, $w(A) = E[\|g\|_\infty]$.

Below we list some simple and useful properties of the Gaussian width of $A \subseteq \mathbb{R}^p$:

Property 1: $w(A) \leq w(B)$ for $A \subseteq B$.

Property 2: $w(A) = w(\text{conv}(A))$, where $\text{conv}(\cdot)$ denotes the convex hull of A.

Property 3: $w(cA) = cw(A)$ for any positive scalar c, in which $cA = \{cx \mid x \in A\}$.

Property 4: $w(\Gamma A) = w(A)$ for any orthogonal matrix $\Gamma \in \mathbb{R}^{p\times p}$.

Property 5: $w(A + b) = w(A)$ for any $A \subseteq \mathbb{R}^p$ and fixed $b \in \mathbb{R}^p$.
The last two properties illustrate the Gaussian width is rotation and translation invariant.

A.2 Sub-Gaussian and Sub-exponential Random Variables (Vectors)

In the proof, we will also frequently use the properties of sub-Gaussian and sub-exponential random variables (vectors). In particular, we are interested in their definitions using moments.

Definition 2 Sub-Gaussian (sub-exponential) random variable: We say that a random variable x is sub-Gaussian (sub-exponential) if the moments satisfies

$$[E|x|^p]^{\frac{1}{p}} \leq K_2\sqrt{p} \quad ([E|x|^p]^{\frac{1}{p}} \leq K_1 p) \tag{69}$$

for any $p \geq 1$ with a constant K_2 (K_1). The minimum value of K_2 (K_1) is called sub-Gaussian (sub-exponential) norm of x, denoted by $\|x\|_{\psi_2}$ ($\|x\|_{\psi_1}$).

Definition 3 Sub-Gaussian (sub-exponential) random vector: We say that a random vector X in $\mathbb{R}^n$ is sub-Gaussian (sub-exponential) if the one-dimensional marginals $\langle X, x\rangle$ are sub-Gaussian (sub-exponential) random variables for all $x \in \mathbb{R}^n$. The sub-Gaussian (sub-exponential) norm of X is defined as

$$\|X\|_{\psi_2} = \sup_{x\in S^{n-1}} \|\langle X, x\rangle\|_{\psi_2} \quad (\|X\|_{\psi_1} = \sup_{x\in S^{n-1}} \|\langle X, x\rangle\|_{\psi_1}) \tag{70}$$

The following definitions and lemmas are from [41].

Lemma 4 *Consider a finite number of independent centered sub-Gaussian random variables X_i. Then $\sum_i X_i$ is also a centered sub-Gaussian random variable. Moreover,*

$$\left\|\sum_i X_i\right\|_{\psi_2}^2 \leq C\sum_i \|X_i\|_{\psi_2}^2 \tag{71}$$

Lemma 5 *Let $X_1, \ldots, X_n$ be independent centered sub-Gaussian random variables. Then $X = (X_1, \ldots, X_n)$ is a centered sub-Gaussian random vector in $\mathbb{R}^n$, and*

$$\|X\|_{\psi_2} \leq C \max_{i\leq n} \|X_i\|_{\psi_2} \tag{72}$$

where C is an absolute constant.

Lemma 6 *Consider a sub-Gaussian random vector X with sub-Gaussian norm $K = \max_i \|X_i\|_{\psi_2}$, then, $Z = \langle X, a\rangle$ is a sub-Gaussian random variable with sub-Gaussian norm $\|Z\|_{\psi_2} \leq CK\|a\|_2$.*

Lemma 7 *A random variable X is sub-Gaussian if and only if X^2 is sub-exponential. Moreover,*

$$|||X|||_{\psi_2}^2 \leq \left|\left|\left|X^2\right|\right|\right|_{\psi_1} \leq 2|||X|||_{\psi_2}^2 \tag{73}$$

Lemma 8 *If X is sub-Gaussian (or sub-exponential), then so is $X - EX$. Moreover, the following holds,*

$$|||X - EX|||_{\psi_2} \leq 2|||X|||_{\psi_2}, \quad |||X - EX|||_{\psi_1} \leq 2|||X|||_{\psi_1} \tag{74}$$

B Restricted Error Set and Recovery Guarantees

Section 2 is about the restricted error set. Lemma 1 characterizes the restricted error set. Theorem 1 establishes the relation between the constrained and restricted error sets. In particular, we prove that the Gaussian width of the regularized and constrained error sets (cone) are of the same order. Starting with the assumption that the RSC condition is satisfied Lemma 2 and Theorem 2 derive results on the upper bound on the L_2 norm of the error.

We collect the proofs of the different results in this section.

B.1 The Restricted Error Set

Lemma 1 in Section 2 characterizes the set to which the error vector belongs. We give the proof of Lemma 1 below:

Lemma 1 *For any $\beta > 1$, assuming*

$$\lambda_n \geq \beta R^*(\nabla \mathcal{L}(\theta^*; Z^n)) \tag{75}$$

where $R^*(\cdot)$ is the dual norm of $R(\cdot)$. Then the error vector $\hat{\Delta}_n = \hat{\theta}_{\lambda_n} - \theta^*$ belongs to the set:

$$E_r = E_r(\theta^*, \beta) = \left\{ \Delta \in \mathbb{R}^p \;\middle|\; R(\theta^* + \Delta) \leq R(\theta^*) + \frac{1}{\beta} R(\Delta) \right\} . \tag{76}$$

Proof: By the optimality of $\hat{\theta}_{\lambda_n} = \theta^* + \hat{\Delta}_n$, we have

$$\mathcal{L}(\theta^* + \hat{\Delta}_n) + \lambda_n R(\theta^* + \hat{\Delta}_n) - \{\mathcal{L}(\theta^*) + \lambda_n R(\theta^*)\} \leq 0 . \tag{77}$$

Now, since $\mathcal{L}$ is convex,

$$\mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) \geq \langle \nabla \mathcal{L}(\theta^*), \Delta \rangle \geq -|\langle \nabla \mathcal{L}(\theta^*), \Delta \rangle| . \tag{78}$$

Further, by generalized Holder's inequality, we have

$$|\langle \nabla \mathcal{L}(\theta^*), \Delta \rangle| \leq R^*(\nabla \mathcal{L}(\theta^*)) R(\Delta) \leq \frac{\lambda_n}{\beta} R(\Delta) , \tag{79}$$

where we have used $\lambda_n \geq \beta R^*(\nabla \mathcal{L}(\theta^*; Z^n))$. Hence, we have

$$\mathcal{L}(\theta^* + \hat{\Delta}_n) - \mathcal{L}(\theta^*) \geq -\frac{\lambda_n}{\beta} R(\hat{\Delta}_n) . \tag{80}$$

As a result,

$$\lambda_n \left\{ R(\theta^* + \hat{\Delta}_n) - R(\theta^*) - \frac{1}{\beta} R(\hat{\Delta}_n) \right\} \leq 0 . \tag{81}$$

Noting that $\lambda_n > 0$ and rearranging completes the proof. ∎

B.2 Relation between the Constrained and Regularized Error Cones

In this section we show that the sizes of the regularized and constrained error sets are of the same order. Recall from [14], that the error set for the constrained setting for atomic norms is a cone given by:

$$C_c = C_c(\theta^*) = \text{cone}(E_c) = \text{cone}\left\{\Delta \in \mathbb{R}^p \mid R(\theta^* + \Delta) \leq R(\theta^*)\right\} . \tag{82}$$

The error set E_r is given by:

$$E_r = E_r(\theta^*, \beta) = \left\{\Delta \in \mathbb{R}^p \;\middle|\; R(\theta^* + \Delta) \leq R(\theta^*) + \frac{1}{\beta} R(\Delta)\right\} .$$

Below we provide the proof of Theorem 1.

Theorem 1 *Let $A_c^{(\rho)} = E_c \cap \rho B_2^p$, $A_r^{(\rho)} = E_r \cap \rho B_2^p$, and $\bar{A}_c^{(\rho)} = C_c \cap \rho B_2^p$, where $\rho B_2^p = \{u \| \|u\|_2 \leq \rho\}$ is the L_2 ball of any radius $\rho > 0$. Then, for any $\beta > 1$ we have*

$$w(A_c^{(\rho)}) \leq w(A_r^{(\rho)}) \leq \left(1 + \frac{2}{\beta - 1} \frac{\|\theta^*\|_2}{\rho}\right) w(\bar{A}_c^{(\rho)}) , \tag{83}$$

where $w(A)$ denotes the Gaussian width of any set A given by: $w(A) = E_g\left[\sup_{a \in A} \langle a, g\rangle\right]$, where g is an isotropic Gaussian random vector, i.e., $g \sim N(0, \mathbb{I}_{p \times p})$.

Proof: The first inequality simply follows from the fact that $E_c \subseteq E_r$ and Property 1 of Gaussian width. For the second part, from triangle inequality, we have

$$R(\Delta) \leq R(\theta^* + \Delta) + R(\theta^*) . \tag{84}$$

Then,

$$\begin{aligned}
E_r(\theta^*, \beta) &= \left\{\Delta \in \mathbb{R}^p \,\middle|\, R(\theta^* + \Delta) \leq R(\theta^*) + \frac{1}{\beta} R(\Delta)\right\} \\
&\subseteq \left\{\Delta \in \mathbb{R}^p \,\middle|\, R(\theta^* + \Delta) \leq R(\theta^*) + \frac{1}{\beta} R(\theta^* + \Delta) + \frac{1}{\beta} R(\theta^*)\right\} \\
&= \left\{\Delta \in \mathbb{R}^p \,\middle|\, \left(1 - \frac{1}{\beta}\right) R(\theta^* + \Delta) \leq \left(1 + \frac{1}{\beta}\right) R(\theta^*)\right\} \\
&= \left\{\Delta \in \mathbb{R}^p \,\middle|\, R(\theta^* + \Delta) \leq \frac{\beta + 1}{\beta - 1} R(\theta^*)\right\} = \tilde{E}_r(\theta^*, \beta) .
\end{aligned}$$

Let $\tilde{C}_r(\theta^*, \beta)$ denote the following set

$$\tilde{C}_r = \tilde{C}_r(\theta^*, \beta) = \text{cone}\left\{\Delta - \frac{2}{\beta - 1}\theta^* \,\middle|\, \Delta \in \bar{E}_r\right\} + \frac{2}{\beta - 1}\theta^* . \tag{85}$$

It follows naturally from the construction that $E_r \subseteq \tilde{C}_r$.

Let $\tilde{A}_r^{(\rho)} = \tilde{C}_r(\theta^*, \beta) \cap \rho B_2^p$. Since $E_r(\theta^*, \beta) \subseteq \tilde{C}_r(\theta^*, \beta)$, we have $w(A_r^{(\rho)}) \leq w(\tilde{A}_r^{(\rho)})$. We define two additional sets for our analysis:

$$\tilde{B}_r^{(\rho)} = \tilde{A}_r^{(\rho)} - \frac{2}{\beta - 1}\theta^* = \left\{\Delta \in \mathbb{R}^p \,\middle|\, \Delta + \frac{2}{\beta - 1}\theta^* \in \tilde{A}_r^{(\rho)}\right\} , \tag{86}$$

$$\tilde{D}_c^{(\rho)} = C_c(\theta^*, \beta) \cap \left(\rho + \frac{2}{\beta - 1}\|\theta^*\|_2\right) B_2^p = \left\{\Delta \in \mathbb{R}^p \,\middle|\, \Delta \in C_c, \|\Delta\|_2 \leq \left(\rho + \frac{2}{\beta - 1}\right)\|\theta^*\|_2\right\} . \tag{87}$$

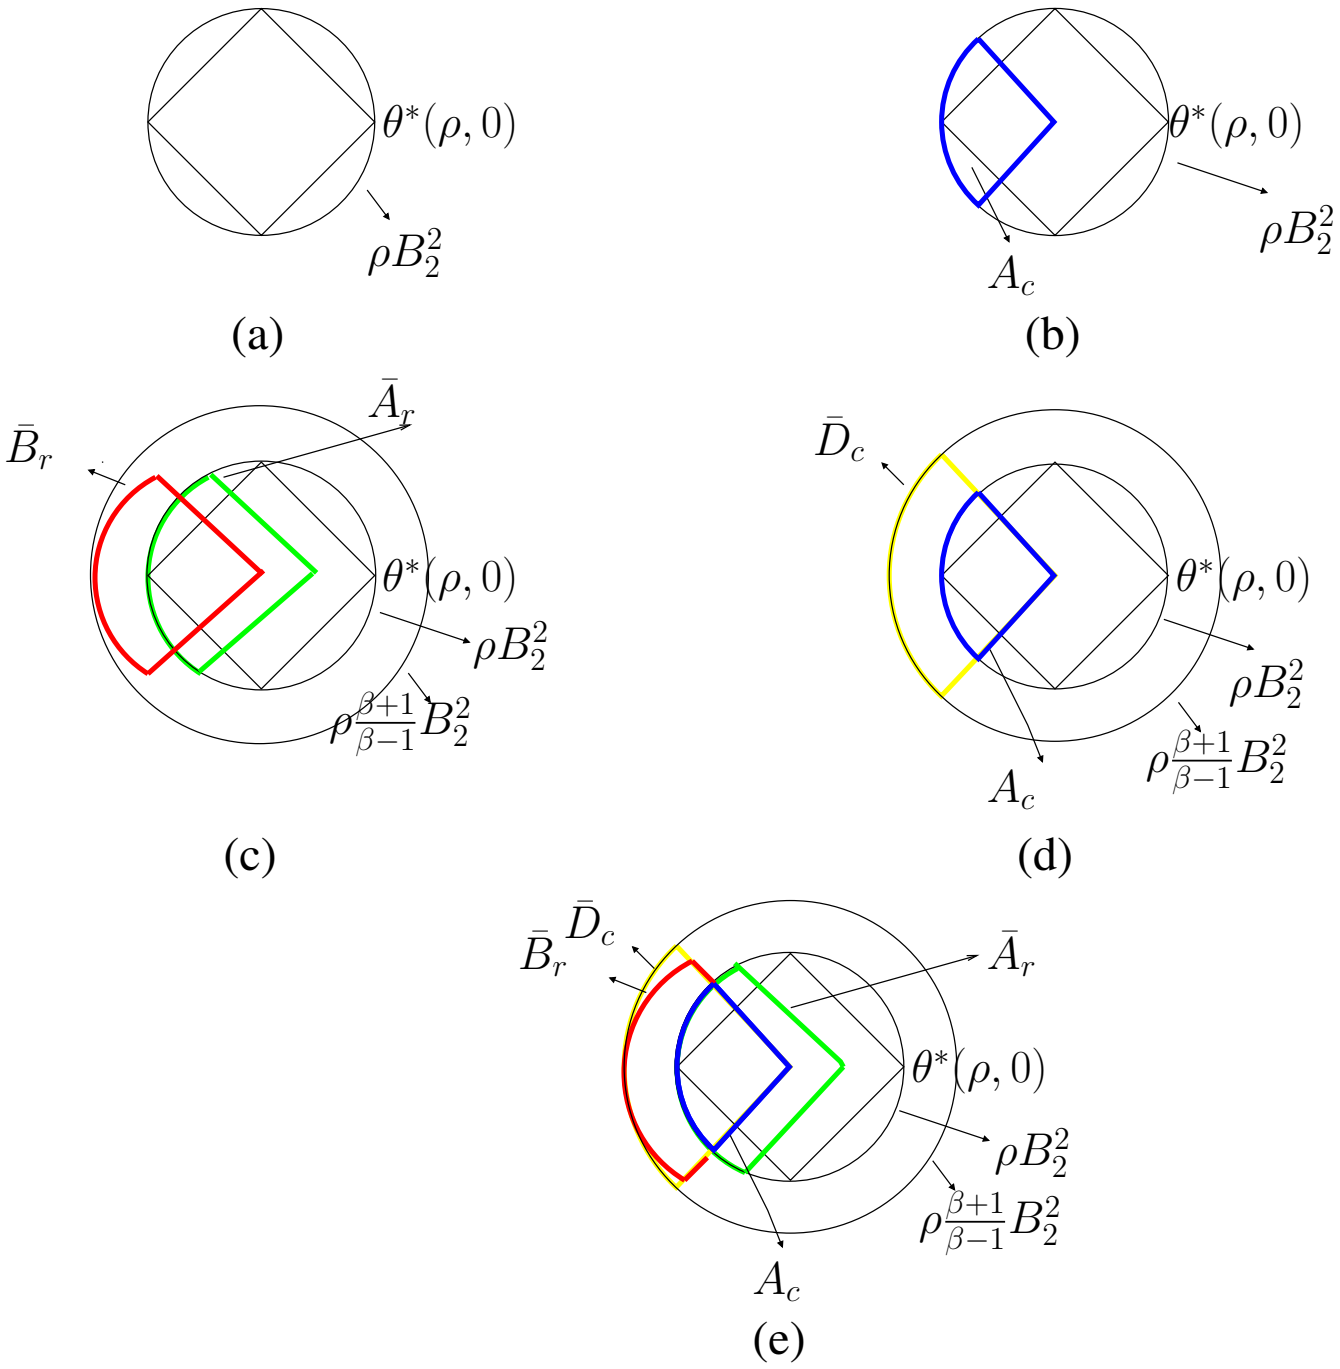

Figure 5: Error cone for L_1 norm in two dimensions: (a) The L1 norm ball in two dimensions; (b) The constrained error cone A_c; (c) The regularized error cone $\hat{A}_r$ and the shifted cone $\hat{B}_r$; (d) The constrained error cone A_c and the shifted constrained error cone $\bar{D}_c$; (e) All error cones.

Following Property 3 of Gaussian width, we have

$$w(\tilde{D}_c^{(\rho)}) = \left(1 + \frac{2}{\beta - 1}\frac{\|\theta^*\|_2}{\rho}\right) w(\bar{A}_c^{(\rho)}) \; . \tag{88}$$

Further, using Property 5 of Gaussian width, we have

$$w(\tilde{A}_r^{(\rho)}) = w(\tilde{B}_r^{(\rho)}) \; . \tag{89}$$

From the construction it is clear that $\tilde{B}_r^{(\rho)} \subset \tilde{D}_c^{(\rho)}$. Hence we have

$$w(\tilde{D}_c^{(\rho)}) \geq w(\tilde{B}_r^{(\rho)}) \tag{90}$$

Then, we have

$$w(\tilde{A}_r^{(\rho)}) = w(\tilde{B}_r^{(\rho)}) \leq w(\tilde{D}_c^{(\rho)}) = \left(1 + \frac{2}{\beta - 1}\frac{\|\theta^*\|_2}{\rho}\right) w(\bar{A}_c^{(\rho)})$$

By noting that $w(A_r^{(\rho)}) \leq w(\tilde{A}_r^{(\rho)})$, we complete the proof. ∎

B.3 Recovery Guarantees

Lemma 2 and Theorem 2 in the paper are results which establish recovery guarantees. The result in Lemma 2 depends on θ^*, which is unknown. On the other hand Theorem 2 gives the result in terms of quantities like λ_n and the norm compatibility constant $\Psi(E_r) = \sup_{u \in E_r} \frac{R(u)}{\|u\|_2}$ which are easier to compute or bound. In this section we give proofs of Lemma 2 and Theorem 2.

Lemma 2 *Assume that the RSC condition is satisfied in E_r by the loss $\mathcal{L}(\cdot)$ with parameter κ. With $\hat{\Delta}_n = \hat{\theta}_{\lambda_n} - \theta^*$, for any norm $R(\cdot)$, we have*

$$\|\hat{\Delta}_n\|_2 \leq \frac{1}{\kappa}\|\nabla\mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*)\|_2 \; , \tag{91}$$

where $\nabla R(\cdot)$ is any sub-gradient of the norm $R(\cdot)$.

Proof: By the RSC property in E_r, for any $\Delta \in E_r$ we have

$$\mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) \geq \langle \nabla L(\theta^*), \Delta \rangle + \kappa \|\Delta\|_2^2 \; . \tag{92}$$

Also, recall that any norm is convex, since by triangle inequality, for $t \in [0, 1]$, we have

$$R(t\theta_1 + (1-t)\theta_2) \leq R(t\theta_1) + R((1-t)\theta_2) = tR(\theta_1) + (1-t)R(\theta_2) \; . \tag{93}$$

As a result, for any sub-gradient $\nabla R(\theta)$ of $R(\theta)$, we have

$$R(\theta^* + \Delta) - R(\theta^*) \geq \langle \Delta, \nabla R(\theta^*) \rangle \; . \tag{94}$$

Adding (92) and (94), we get

$$\mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) + \lambda_n(R(\theta^* + \Delta) - R(\theta^*)) \geq \langle \nabla\mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*), \Delta \rangle + \kappa \|\Delta\|_2^2 \tag{95}$$

Now, by Cauchy-Schwartz inequality, we have

$$\begin{aligned} & |\langle \nabla \mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*), \Delta \rangle| \leq \|\nabla \mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*)\|_2 \|\Delta\|_2 \\ \Rightarrow \quad & \langle \nabla \mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*), \Delta \rangle \geq -\|\nabla \mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*)\|_2 \|\Delta\|_2 \, . \end{aligned} \tag{96}$$

Using (96) in (95), we have

$$\begin{aligned} \mathcal{F}(\Delta) &= \mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) + \lambda_n (R(\theta^* + \Delta) - R(\theta^*)) \\ &\geq -\|\nabla \mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*)\|_2 \|\Delta\|_2 + \kappa \|\Delta\|_2^2 \\ &= \kappa \|\Delta\|_2 \left\{ \|\Delta\|_2 - \frac{\|\nabla \mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*)\|_2}{\kappa} \right\} \, . \end{aligned} \tag{97}$$

Now, since $\mathcal{F}(\hat{\Delta}_n) \leq 0$, from (97), we have

$$\|\hat{\Delta}_n\|_2 \leq \frac{\|\nabla \mathcal{L}(\theta^*) + \lambda_n \nabla R(\theta^*)\|_2}{\kappa} \, , \tag{98}$$

which completes the proof. ∎

Theorem 2 *Assume that the RSC condition is satisfied in E_r by the loss $\mathcal{L}(\cdot)$ with parameter κ. With $\hat{\Delta}_n = \hat{\theta}_{\lambda_n} - \theta^*$, for any norm $R(\cdot)$, we have*

$$\|\hat{\Delta}_n\|_2 \leq \frac{1+\beta}{\beta} \frac{\lambda_n}{\kappa} \Psi(E_r) \, . \tag{99}$$

Proof: By the RSC property in E_r, we have for any $\Delta \in E_r$

$$\mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) \geq \langle \nabla L(\theta^*), \Delta \rangle + \kappa \|\Delta\|_2^2 \, . \tag{100}$$

By definition of a dual norm, we have

$$|\langle \nabla \mathcal{L}(\theta^*), \Delta \rangle| \leq R^*(\nabla \mathcal{L}(\theta^*)) R(\Delta) \, . \tag{101}$$

Further, by construction, $R^*(\nabla \mathcal{L}(\theta^*)) \leq \frac{\lambda_n}{\beta}$, implying

$$\begin{aligned} & |\langle \nabla \mathcal{L}(\theta^*), \Delta \rangle| \leq \frac{\lambda_n}{\beta} R(\Delta) \\ \Rightarrow \quad & \langle \nabla \mathcal{L}(\theta^*), \Delta \rangle \geq -\frac{\lambda_n}{\beta} R(\Delta) \, . \end{aligned} \tag{102}$$

Further, from triangle inequality, we have

$$R(\theta^* + \Delta) - R(\theta^*) \geq -R(\Delta) \tag{103}$$

Adding (102) and (103), we have

$$\begin{aligned} \mathcal{F}(\Delta) = \mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) + \lambda_n (R(\theta^* + \Delta) - R(\theta^*)) &\geq -\frac{\lambda_n}{\beta} R(\Delta) + \kappa \|\Delta\|_2^2 - \lambda_n R(\Delta) \\ = \kappa \|\Delta\|_2^2 - \lambda_n \frac{1+\beta}{\beta} R(\Delta) \, . \end{aligned} \tag{104}$$

By definition of the norm compatibility constant ψ_{E_r}, we have $R(\Delta) \leq \|\Delta\|_2 \psi_{E_r}$ implying $-R(\Delta) \geq -\|\Delta\|_2 \psi_{E_r}$. Plugging the inequality back into (104), we have

$$\mathcal{F}(\Delta) \geq \kappa \|\Delta\|_2 \left\{ \|\Delta\|_2 - \frac{1+\beta}{\beta} \frac{\lambda_n}{\kappa} \psi_{E_r} \right\} . \tag{105}$$

Since $\mathcal{F}(\hat{\Delta}_n) \leq 0$, we have

$$\|\hat{\Delta}_n\|_2 \leq \frac{1+\beta}{\beta} \frac{\lambda_n}{\kappa} \psi_{E_r} , \tag{106}$$

which completes the proof. ∎

C Bounds on the Regularization Parameter

In this section, we prove Theorem 3 and 4 in Section 3 of the paper. The regularization parameter should satisfy the condition $\lambda_n \geq \beta R^*(\nabla \mathcal{L}(\theta^*; Z^n))$. In Theorem 3 we establish the upper bound on the expectation $E[R^*(\nabla \mathcal{L}(\theta^*; Z^n))]$ in terms of the Gaussian width of the unit norm ball for least squares loss and Gaussian designs. In Theorem 4 we show that $R^*(\nabla \mathcal{L}(\theta^*; Z^n))$ concentrates sharply around its expectation.

C.1 Proof of Theorem 3

To prove Theorem 3, we first need the following theorem from generic chaining.

Theorem 8 *Let $\Omega_R = \{u : R(u) \leq 1\}$ be the unit norm ball of $R(\cdot)$. Assuming h is any centered sub-Gaussian random vector with $\|\|h\|\|_{\psi_2} \leq \kappa$, then we have*

$$E\left[\sup_{R(u) \leq 1} \langle h, u \rangle \right] \leq \eta_0 \kappa w(\Omega_R) ,$$

where η_0 is a universal constant.

Proof: The quantity $E[\sup_{R(u)\leq 1} \langle h, u \rangle]$ can be considered the "sub-Gaussian width" of Ω_R, the unit norm ball, since it has the exact same form as the Gaussian width, with h being a sub-Gaussian vector instead of a Gaussian vector. Next, we show that the sub-Gaussian width is always bounded by the Gaussian width times a factor proportional to κ.

Consider the sub-Gaussian process $Y = \{Y_u\}, Y_u = \langle u, h \rangle$ indexed by $u \in \Omega_R$, the unit norm ball. Consider the Gaussian process $X = \{X_u\}, X_u = \langle u, g \rangle$, where $g \sim N(0, \mathbb{I})$, indexed by the same set, i.e., $u \in \Omega_R$, the unit norm ball. First, note that $|Y_u - Y_v| = |\langle h, u - v \rangle|$, so that by the concentration of sub-Gaussian random variable [41, Equation 5.10], we have

$$P\left(|Y_u - Y_v| \geq \epsilon \right) \leq e \cdot \exp\left(-\frac{c\epsilon^2}{\kappa^2 \|u - v\|^2} \right) , \tag{107}$$

where $c > 0$ is an absolute constant. As a result, a direct application of the generic chaining argument for upper bounds on such empirical processes [37, Theorem 2.1.5] gives

$$E\left[\sup_{u,v} |Y_u - Y_v| \right] \leq \eta_1 E\left[\sup_u X_u \right] = \eta_1 w(\Omega_R) , \tag{108}$$

where η_1 is an absolute constant. Further, since $\{Y_u\}$ is a symmetric process, from [37, Lemma 1.2.8], we have

$$E\left[\sup_{u,v}|Y_u - Y_v|\right] = 2E\left[\sup_u Y_u\right] . \tag{109}$$

As a result, with $\eta_0 = \eta_1/2$, we have

$$E\left[\sup_{R(u)\le 1}\langle h, u\rangle\right] = E\left[\sup_u Y_u\right] \le \eta_0 w(\Omega_R) . \tag{110}$$

That completes the proof. ∎

Now we turn to the proof of Theorem 3.

Theorem 3 *Let $\Omega_R = \{u : R(u) \le 1\}$, and $\mathcal{L}$ be the squared loss. For sub-Gaussian design X and noise ω, we have*

$$E\left[R^*(\nabla\mathcal{L}(\theta^*; Z^n))\right] \le \eta\xi \cdot \frac{\kappa w(\Omega_R)}{\sqrt{n}} , \tag{111}$$

where the expectation is taken over both X and ω. The constant ξ is given by

$$\xi = \begin{cases} 1 & \text{if } X \text{ is isotropic} \\ \sqrt{\Lambda_{\max}(\Sigma)} & \text{if } X \text{ is anisotropic} . \end{cases}$$

Proof: For least squares loss, we first note that

$$\begin{aligned} E\left[R^*(\nabla\mathcal{L}(\theta^*; Z^n))\right] &= E\left[R^*(\frac{1}{n}X^T\omega)\right] = E\left[\sup_{R(u)\le 1}\langle\frac{1}{n}X^T\omega, u\rangle\right] \\ &= E\left[\frac{1}{n}\|\omega\|_2 \cdot E\left[\sup_{R(u)\le 1}\left\langle X^T\frac{\omega}{\|\omega\|_2}, u\right\rangle\Big|\omega\right]\right] \\ &\le E\left[\frac{1}{n}\|\omega\|_2 \cdot \sup_{v\in\mathbb{S}^{n-1}} E\left[\sup_{R(u)\le 1}\left\langle X^T v, u\right\rangle\right]\right] \\ &= \frac{1}{n}E\left[\|\omega\|_2\right] \cdot \sup_{v\in\mathbb{S}^{n-1}} E\left[\sup_{R(u)\le 1}\left\langle X^T v, u\right\rangle\right] . \end{aligned} \tag{112}$$

$E[\|\omega\|_2]$ is the expected length of a centered sub-Gaussian random vector, which can be easily bounded using Jensen's inequality,

$$E[\|\omega\|_2] < \sqrt{E[\|\omega\|_2^2]} = \sqrt{n} . \tag{113}$$

Then we focus on $E\left[\sup_{R(u)\le 1}\langle X^T v, u\rangle\right]$ for any fixed $v \in \mathbb{S}^{n-1}$. Let $h = X^T v$ be a random vector, and consider the random variable $\langle h, z\rangle$ for any fixed $z \in \mathbb{S}^{p-1}$. Note that

$$\langle h, z\rangle = \langle v, Xz\rangle .$$

Case 1. If X is independent isotropic, then Xz has i.i.d. centered sub-Gaussian entries with ψ_2 norm at most κ. By Lemma 4, we know that $\langle v, Xz\rangle$ is sub-Gaussian with $\|\!|\langle v, Xz\rangle|\!\|_{\psi_2} \le C\kappa$, where C is an absolute constant. Hence h is a sub-Gaussian random vector with $\|\!|h|\!\|_{\psi_2} \le C\kappa$. Using Theorem 8, we conclude that for any $v \in \mathbb{S}^{n-1}$

$$E\left[\sup_{R(u)\le 1}\left\langle X^T v, u\right\rangle\right] \le \eta_0 C\kappa \cdot w(\Omega_R) . \tag{114}$$

Case 2. If X is independent anisotropic, then Xz has i.i.d. centered sub-Gaussian entries with ψ_2 norm at most $\kappa\sqrt{\Lambda_{\max}(\Sigma)}$, where $\Lambda_{\max}(\Sigma)$ is the largest eigenvalue of Σ. By the same argument as Case 1, we have

$$E\left[\sup_{R(u)\leq 1}\left\langle X^T v, u\right\rangle\right] \leq \eta_0 C\kappa\sqrt{\Lambda_{\max}(\Sigma)}\cdot w(\Omega_R)\ . \tag{115}$$

Letting $\eta = \eta_0 C$ and combining (114), (113) and (115), we complete the proof. ∎

C.2 Proof of Theorem 4

To prove Theorem 4, we also need the following result from generic chaining.

Theorem 9 *Let $\Omega_R = \{u : R(u) \leq 1\}$ be the unit norm ball of $R(\cdot)$. Assuming h is any centered sub-Gaussian random vector with $\|\|h\|\|_{\psi_2} \leq \kappa$, then we have for any $\tau > 0$,*

$$P\left(\sup_{R(u)\leq 1}\langle h, u\rangle \geq \nu_0\kappa w(\Omega_R) + \tau\right) \leq \nu_1 \exp\left(-\left(\frac{\tau}{\nu_2\kappa\phi}\right)^2\right)\ , \tag{116}$$

where ν_0, ν_1 and ν_2 are universal constants, and $\phi = \sup_{R(u)\leq 1}\|u\|_2$.

Proof: Consider the sub-Gaussian process $Y = \{Y_u\}, Y_u = \langle h, u\rangle$ indexed by $u \in \Omega_R$. By the same argument in the proof of Theorem 8, we have

$$P\left(|Y_u - Y_v| \geq \epsilon\right) \leq e\cdot\exp\left(-\frac{c\epsilon^2}{\kappa^2\|u-v\|^2}\right)\ , \tag{117}$$

and

$$E\left[\sup_{u,v}|Y_u - Y_v|\right] = 2E\left[\sup_u Y_u\right]\ . \tag{118}$$

Then a direct application of [37, Theorem 2.1.5] and [38, Theorem 2.2.27] gives us (116). ∎

Theorem 4 *Let design X and noise ω be sub-Gaussian, and $\mathcal{L}$ be squared loss. Define $\phi = \sup_{R(u)\leq 1}\|u\|_2$, then for any $\tau > 0$, with probability at least $1 - c_1\exp\left(-\min\left((\frac{\tau}{c_2\xi\kappa\phi})^2, c_0 n\right)\right)$, we have*

$$R^*\left(\nabla\mathcal{L}(\theta^*; Z^n)\right) \leq \sqrt{\frac{2K^2+1}{n}}\left(c\xi\kappa\cdot w(\Omega_R) + \tau\right)\ , \tag{119}$$

where c, c_0, c_1 and c_2 are all absolute constants, and ξ is the same as in Theorem 3.
Proof: We only show the case for isotropic X, where $\xi = 1$. Note that

$$\begin{aligned}
&P\left(R^*\left(\nabla\mathcal{L}(\theta^*; Z^n)\right) \geq \sqrt{\frac{2K^2+1}{n}}\left(c\xi\kappa\cdot w(\Omega_R) + \tau\right)\right)\\
&= P\left(\|\omega\|_2\cdot R^*\left(\frac{X^T\omega}{\|\omega\|_2}\right) \geq \sqrt{(2K^2+1)n}\left(c\kappa\cdot w(\Omega_R) + \tau\right)\right)\\
&\leq P\left(\|\omega\|_2 > \sqrt{(2K^2+1)n}\right) + \sup_{v\in\mathbb{S}^{n-1}} P\left(R^*\left(X^T v\right) \geq c\kappa\cdot w(\Omega_R) + \tau\right)\ ,
\end{aligned} \tag{120}$$

where the last inequality uses the union bound. We first prove the bound for $\|\omega\|_2$. Since ω consists of i.i.d. centered unit-variance sub-Gaussian elements with $\|\omega_i\|_{\psi_2} < K$, ω_i^2 is sub-exponential with $\|\omega_i\|_{\psi_1} < 2K^2$. By applying Bernstein's inequality to $\|\omega\|_2^2 = \sum_{i=1}^n \omega_i^2$, we obtain

$$P\left(\left|\|\omega\|_2^2 - E[\|\omega\|_2^2]\right| \geq \tau\right) \leq 2\exp\left[-c_0 \min\left(\frac{\tau^2}{4K^4 n}, \frac{\tau}{2K^2}\right)\right] ,$$

where c_0 is an absolute constant. Setting $\tau = 2K^2 n$ and using (113), we get

$$P\left(\|\omega\|_2 \geq \sqrt{(2K^2+1)n}\right) \leq 2\exp(-c_0 n) \tag{121}$$

Next we bound $R^*(X^T v)$ for any $v \in \mathbb{S}^{n-1}$. Given any fixed $v \in \mathbb{S}^{n-1}$, we note that

$$R^*(X^T v) = \sup_{R(u)\leq 1} \langle X^T v, u\rangle ,$$

and $X^T v$ is a sub-Gaussian random vector with $\left\|\left\|X^T v\right\|\right\|_{\psi_2} \leq C\kappa$ as shown in the proof of Theorem 3. Using Theorem 9, we have

$$P\left(R^*(X^T v) \geq \nu_0 C\kappa \cdot w(\Omega_R) + \tau\right) \leq \nu_1 \exp\left(-\left(\frac{\tau}{\nu_2 C\kappa\phi}\right)^2\right) . \tag{122}$$

Letting $c = \nu_0 C$, $c_1 = \nu_1 + 2$ and $c_2 = \nu_2 C$, and combining (121) and (122), we complete the proof. ∎

C.3 Proof of Lemma 3

Lemma 3 *Let $\Omega_R = \{u : R(u) \leq 1\}$ be the unit norm ball and $\Theta_R = \{u : R(u) = 1\}$ be the boundary. For any $\tilde{\theta} \in \Theta_R$ define $\rho(\tilde{\theta}) = \sup_{\theta: R(\theta)\leq 1} \|\theta - \tilde{\theta}\|_2$ is the diameter of Ω_R measured with respect to $\tilde{\theta}$. If $G(\tilde{\theta}) = \mathrm{cone}(\Omega_R - \tilde{\theta}) \cap \rho(\tilde{\theta}) B_2^p$, i.e., the cone of $(\Omega_R - \tilde{\theta})$ intersecting the ball of radius $\rho(\tilde{\theta})$. Then*

$$w(\Omega_R) \leq \inf_{\tilde{\theta}\in\Theta_R} w(G(\tilde{\theta})) \tag{123}$$

Proof: For any $\tilde{\theta} \in \Theta_R$, consider the set $F_R(\tilde{\theta}) = \Omega_R - \tilde{\theta} = \{u : R(u + \tilde{\theta}) \leq 1\}$. Since Gaussian width is translation invariant, the Gaussian width of Ω_R and F_R are the same, i.e., $w(\Omega_R) = w(F_R(\tilde{\theta}))$. Since, $\rho(\tilde{\theta}) = \sup_{\theta: R(\theta)\leq 1} \|\theta - \tilde{\theta}\|_2$ is the diameter of Ω_R as well as $F_R(\tilde{\theta})$, a ball of radius $\rho(\tilde{\theta})$ will include $F_R(\tilde{\theta})$, so that $F_R(\tilde{\theta}) \subseteq \rho(\tilde{\theta}) B_2^p$. Further, by definition, $F_R(\tilde{\theta}) \subseteq \mathrm{cone}(F_R(\tilde{\theta})) = \mathrm{cone}(\Omega_R - \tilde{\theta})$. Let $G(\tilde{\theta}) = \mathrm{cone}(\Omega_R - \tilde{\theta}) \cap \rho(\tilde{\theta}) B_2^p$. By construction, $F_R(\tilde{\theta}) \subseteq G(\tilde{\theta})$. Then,

$$w(\Omega_R) = w(F_R(\tilde{\theta})) \leq w(G(\tilde{\theta})) .$$

Noting the analysis holds for any $\tilde{\theta} \in \Theta_R$, completes the proof. ∎

D Restricted Eigenvalue Conditions: Sub-Gaussian Designs

We focus on results in Section 4. In particular we consider RE conditions for sub-Gaussian design matrices for three different cases: (i) the design matrix has independent sub-Gaussian rows X_i with $\|X_i\|_{\psi_2} \leq \kappa$, (ii) the design matrix has independent rows with subGaussian elements x_{ij} so that $\|x_{ij}\|_{\psi_2} \leq \kappa$ and the

columns are correlated, and (iii) the columns are independent subGaussian but the rows are correlated, i.e., correlated samples. One can view (ii) as a special case of (i), but we highlight this special case because of its practical importance and past literature on RE conditions for anisotropic subGaussian designs [33, 34].

Our results simply use a general treatment developed in [28], building on [21], based on Talagrand's generic chaining [37, 38]. We specifically focus on results in [28] which provide uniform bounds on the supremum of certain empirical processes. RE results for the specific cases of interest in the current paper will then be established by suitable choices of these empirical processes. The results in [28], and more generally in generic chaining [37, 38], are based on certain γ-functionals which we briefly introduce below.

Consider a metric space (T, d) and for a finite set $\mathcal{A} \subset T$, let $|\mathcal{A}|$ denote its cardinality. An *admissible sequence* is an increasing sequence of subsets $\{\mathcal{A}_n, n \geq 0\}$ of T, such that $|\mathcal{A}_0| = 1$ and for $n \geq 1$, $|\mathcal{A}_n| = 2^{2^n}$. Given $\alpha > 0$, we define the γ_α-functional as

$$\gamma_\alpha(T, d) = \inf \sup_{t \in T} \sum_{n=0}^{\infty} \text{Diam}(A_n(t)) \,, \tag{124}$$

where $A_n(t)$ is the unique element of $\mathcal{A}_n$ that contains t, $\text{Diam}(A_n(t))$ is the diameter of A_n according to d, and the infimum is over all admissible sequences of T. To get the desired RIP results in terms of Gaussian widths, we start with the following key result, originally [28, Theorem D].

Theorem 10 (Mendelson, Pajor, Tomczak-Jaegermann [28]) *There exist absolute constants c_1, c_2, c_3 for which the following holds. Let (Ω, μ) be a probability space, set F be a subset of the unit sphere of $L_2(\mu)$, i.e., $F \subseteq S_{L_2} = \{f : |||f|||_{L_2} = 1\}$, and assume that $\sup_{f \in F} |||f|||_{\psi_2} \leq \kappa$. Then, for any $\theta > 0$ and $n \geq 1$ satisfying*

$$c_1 \kappa \gamma_2(F, |||\cdot|||_{\psi_2}) \leq \theta \sqrt{n} \,, \tag{125}$$

with probability at least $1 - \exp(-c_2 \theta^2 n / \kappa^4)$,

$$\sup_{f \in F} \left| \frac{1}{n} \sum_{i=1}^{n} f^2(X_i) - E\left[f^2\right] \right| \leq \theta \,. \tag{126}$$

Further, if F is symmetric, then

$$E\left[\sup_{f \in F} \left| \frac{1}{n} \sum_{i=1}^{n} f^2(X_i) - E\left[f^2\right] \right| \right] \leq c_3 \max\left\{ 2\kappa \frac{\gamma_2(F, |||\cdot|||_{\psi_2})}{\sqrt{n}}, \frac{\gamma_2^2(F, |||\cdot|||_{\psi_2})}{n} \right\} \tag{127}$$

We use the above result and related arguments to establish RIP conditions for the cases of interest.

D.1 Isotropic Sub-Gaussian Designs

We consider the case where the design matrix $X \in \mathbb{R}^{n \times p}$ has independent subGaussian rows where each row satisfies $|||X_i|||_{\psi_2} \leq \kappa$ and $E[X_i X_i^T] = \mathbb{I}_{p \times p}$. Thus, the measure μ from which the rows X_i are sampled independently is an isotropic sub-Gaussian measure.

Theorem 11 *Let X be a design matrix with independent isotropic subGaussian rows, i.e., $|||X_i|||_{\psi_2} \leq \kappa$ and $E[X_i X_i^T] = \mathbb{I}$. Then, for absolute constants $\eta, c > 0$, with probability at least $(1 - 2\exp(-\eta w^2(A)))$, we have*

$$\sup_{u \in A} \left| \frac{1}{n} ||Xu||^2 - 1 \right| \;=\; \sup_{u \in A} \left| \frac{1}{n} \sum_{i=1}^{n} \langle X_i, u \rangle^2 - 1 \right| \;\leq\; c \frac{w(A)}{\sqrt{n}} \,, \tag{128}$$

or, equivalently,

$$1 - c\frac{w(A)}{\sqrt{n}} \;\leq\; \inf_{u\in A} \frac{1}{n}||Xu||^2 \;\leq\; \sup_{u\in A} \frac{1}{n}||Xu||^2 \;\leq\; 1 + c\frac{w(A)}{\sqrt{n}} \,. \tag{129}$$

Proof: The result essentially follows from an application of Theorem 10. For convenience of notation, let X_0 be i.i.d. as the rows $X_i, i = 1, \ldots, n$, thus distributed following μ. To apply Theorem 10, we choose any $A \subseteq S^{p-1}$ consider the following class of functions: $F = \{\langle \cdot, u\rangle : u \in A\}$. Then, $f(X_0) = \langle X_0, u\rangle$ and F is a subset of the unit sphere, i.e., $F \subseteq S_{L_2}$, since $|||f|||_{L_2} = E[u^T X_0^T X_0 u] = \|u\|^2 = 1$. Further, $\sup_{f\in F} |||f|||_{\psi_2} = \sup_{u\in A} |||\langle X_0, u\rangle|||_{\psi_2} \leq |||X_0|||_{\psi_2} \leq \kappa/2$.

Next, we show that for the current setting, the γ_2-functional can be upper bounded by $w(A)$, the Gaussian width of A. Since μ is isotropic subGaussian with ψ_2-norm bounded by κ, we have

$$\gamma_2(F \cap S_{L_2}, |||\cdot|||_{\psi_2}) \;\leq\; \kappa\gamma_2(F \cap S_{L_2}, |||\cdot|||_{L_2}) \;\leq\; \kappa c_4 w(A) \,, \tag{130}$$

where the last inequality follows from generic chaining, in particular [37, Theorem 2.1.1], for an absolute constant $c_4 > 0$.

In the context of Theorem 10, we choose

$$\theta = c_1 c_4 \kappa^2 \frac{w(A)}{\sqrt{n}} \geq c_1 \kappa \frac{\gamma_2(F \cap S_{L_2}, |||\cdot|||_{\psi_2})}{\sqrt{n}} \,,$$

so that the condition on θ is satisfied. With this choice of θ, we have

$$\theta^2 n/\kappa^4 = c_1^2 c_4^2 w^2(A) \,.$$

Then, from Theorem 10, it follows that with probability at least $1 - \exp(-\eta w^2(A))$, we have

$$\sup_{u\in A} \left| \frac{1}{n}\sum_{i=1}^n \langle X_i, u\rangle^2 - 1 \right| \;\leq c\frac{w(A)}{\sqrt{n}} \,, \tag{131}$$

where $\eta = c_2 c_1^2 c_4^2$ and $c = c_1 c_2 \kappa^2$ are absolute constants. As a result, we have

$$\sup_{u\in A} \left(\frac{1}{n}\sum_{i=1}^n \langle X_i, u\rangle^2 - 1 \right) \;\leq\; c\frac{w(A)}{\sqrt{n}} \,, \qquad \text{and} \qquad \sup_{u\in A} \left(1 - \frac{1}{n}\sum_{i=1}^n \langle X_i, u\rangle^2 \right) \;\leq\; c\frac{w(A)}{\sqrt{n}} \,,$$

yielding

$$1 - c\frac{w(A)}{\sqrt{n}} \;\leq\; \inf_{u\in A} \frac{1}{n}||Xu||^2 \;\leq\; \sup_{u\in A} \frac{1}{n}||Xu||^2 \;\leq\; 1 + c\frac{w(A)}{\sqrt{n}} \,. \tag{132}$$

That completes the proof. ∎

D.2 Anisotropic Sub-Gaussian Designs

We now consider the case where the design matrix $X \in \mathbb{R}^{n\times p}$ has independent rows, and each row X_i is anisotropic subGaussian with $E[X_i^T X_i] = \Sigma$. Further, we assume that corresponding isotropic random vector $\tilde{X}_i = X_i \Sigma^{-1/2}$ satisfies $\left|\left|\left|\tilde{X}_i\right|\right|\right|_{\psi_2} \leq \kappa$. A simple special case of such an anisotropic subGaussian design is when $X_i \sim N(0, \Sigma)$, where $\tilde{X}_i = X_i\Sigma^{-1/2} \sim N(0, \mathbb{I})$ so that $\left|\left|\left|\tilde{X}_i\right|\right|\right|_{\psi_2} = 1$. The result below characterizes RIP-style property of any such anisotropic subGaussian designs.

Theorem 12 *Let X be a design matrix with independent anisotropic subGaussian rows, i.e., $E[X_i^T X_i] = \Sigma$ and $\left|\left|\left| X_i \Sigma^{-1/2} \right|\right|\right|_{\psi_2} \leq \kappa$. Then, for absolute constants $\eta, c > 0$, with probability at least $(1-2\exp(-\eta w^2(A)))$, we have*

$$\sup_{u \in A} \left| \frac{1}{n} \frac{1}{u^T \Sigma u} ||Xu||^2 - 1 \right| \;=\; \sup_{u \in A} \left| \frac{1}{n} \frac{1}{u^T \Sigma u} \sum_{i=1}^n \langle X_i, u \rangle^2 - 1 \right| \;\leq\; c \frac{w(A)}{\sqrt{n}} \; . \tag{133}$$

Further,

$$\lambda_{\min}(\Sigma|A) \left(1 - c \frac{w(A)}{\sqrt{n}} \right) \;\leq\; \inf_{u \in A} \frac{1}{n} ||Xu||^2 \;\leq\; \sup_{u \in A} \frac{1}{n} ||Xu||^2 \;\leq\; \lambda_{\max}(\Sigma|A) \left(1 + c \frac{w(A)}{\sqrt{n}} \right) , \tag{134}$$

where

$$\lambda_{\min}(\Sigma|A) = \inf_{u \in A} \; u^T \Sigma u \; , \qquad \textit{and} \qquad \lambda_{\max}(\Sigma|A) = \sup_{u \in A} \; u^T \Sigma u \tag{135}$$

are the restricted minimum and maximum eigenvalues of Σ restricted to $A \subseteq S^{p-1}$.

Proof: The result also follows from an application of Theorem 10. For convenience of notation, let X_0 be i.i.d. as the rows $X_i, i = 1, \ldots, n$, thus distributed following μ. To apply Theorem 10, we choose any $A \subseteq S^{p-1}$ consider the following class of functions:

$$F = \{ f_u, u \in A : f_u(\cdot) = \frac{1}{\sqrt{u^T \Sigma u}} \langle \cdot, u \rangle \} \; . \tag{136}$$

Then, $f_u(X_0) = \frac{1}{\sqrt{u^T \Sigma u}} \langle X_0, u \rangle$ and F is a subset of the unit sphere, i.e., $F \subseteq S_{L_2}$, since for $f_u \in F$

$$|||f_u|||_{L_2}^2 = \frac{1}{u^T \Sigma u} E[u^T X_0^T X_0 u] = 1 \; .$$

Next, we focus on getting an upper bound on $\sup_{f_u \in F} \; |||f_u|||_{\psi_2} = \sup_{u \in A} \left|\left|\left| \frac{1}{\sqrt{u^T \Sigma u}} \langle X_0, u \rangle \right|\right|\right|_{\psi_2}$. Let $\tilde{X}_0 = X_0 \Sigma^{-1/2}$ so that $\tilde{X}_0$ is a isotropic vector with $\left|\left|\left| \tilde{X}_0 \right|\right|\right|_{\psi_2} \leq \kappa$. Noting that

$$\langle X_0, u \rangle = \langle \tilde{X}_0, \Sigma^{1/2} u \rangle = \sqrt{u^T \Sigma u} \langle \tilde{X}_0, \frac{\Sigma^{1/2} u}{\|\Sigma^{1/2} u\|_2} \rangle \; ,$$

we have

$$\sup_u |||f_u|||_{\psi_2} = \sup_{u \in A} \left|\left|\left| \frac{1}{\sqrt{u^T \Sigma u}} \langle X_0, u \rangle \right|\right|\right|_{\psi_2} = \sup_{u \in A} \left|\left|\left| \langle \tilde{X}_0, \frac{\Sigma^{1/2} u}{\|\Sigma^{1/2} u\|_2} \rangle \right|\right|\right|_{\psi_2} \leq \left|\left|\left| \tilde{X}_0 \right|\right|\right|_{\psi_2} \leq \kappa \; .$$

As a result, we have

$$\gamma_2(F \cap S_{L_2}, |||\cdot|||_{\psi_2}) \;\leq\; \kappa \, \gamma_2(F \cap S_{L_2}, |||\cdot|||_{L_2}) \;\leq\; \kappa c_4 w(A) \; , \tag{137}$$

where the last inequality follows from [37, Theorem 2.1.1], for an absolute constant $c_4 > 0$.
In the context of Theorem 10, we choose

$$\theta = c_1 c_4 \kappa^2 \frac{w(A)}{\sqrt{n}} \geq c_1 \kappa \frac{\gamma_2(F \cap S_{L_2}, |||\cdot|||_{\psi_2})}{\sqrt{n}} \; ,$$

so that the lower bound condition on θ is satisfied. With this choice of θ, we have

$$\theta^2 n / \kappa^4 = c_1^2 c_4^2 w^2(A) \; .$$

Then, from Theorem 10, it follows that with probability at least $1 - \exp(-\eta w^2(A))$, we have

$$\sup_{u \in A} \left| \frac{1}{n} \frac{1}{u^T \Sigma u} \sum_{i=1}^{n} \langle X_i, u \rangle^2 - 1 \right| \leq c \frac{w(A)}{\sqrt{n}} , \tag{138}$$

where $\eta = c_2 c_1^2 c_4^2$ and $c = c_1 c_2 \kappa^2$ are absolute constants. As a result, we have

$$\sup_{u \in A} \left(\frac{1}{n} \frac{1}{u^T \Sigma u} \sum_{i=1}^{n} \langle X_i, u \rangle^2 - 1 \right) \leq c \frac{w(A)}{\sqrt{n}} , \tag{139}$$

$$\sup_{u \in A} \left(1 - \frac{1}{n} \frac{1}{u^T \Sigma u} \sum_{i=1}^{n} \langle X_i, u \rangle^2 \right) \leq c \frac{w(A)}{\sqrt{n}} . \tag{140}$$

From (139), we have

$$\frac{1}{\lambda_{\max}(\Sigma|A)} \sup_{u \in A} \frac{1}{n} \|Xu\|^2 \leq \sup_{u \in A} \frac{1}{n} \frac{1}{u^T \Sigma u} \|Xu\|^2 \leq 1 + c \frac{w(A)}{\sqrt{n}} ,$$

so that

$$\sup_{u \in A} \frac{1}{n} \|Xu\|^2 \leq \lambda_{\max}(\Sigma|A) + c \lambda_{\max}(\Sigma|A) \frac{w(A)}{\sqrt{n}} . \tag{141}$$

Similarly, from (140), we have

$$1 - c \frac{w(A)}{\sqrt{n}} \leq \inf_{u \in A} \frac{1}{n} \frac{1}{u^T \Sigma u} \|Xu\|^2 \leq \frac{1}{\lambda_{\min}(\Sigma|A)} \inf_{u \in A} \frac{1}{n} \|Xu\|^2 ,$$

implying

$$\lambda_{\min}(\Sigma|A) - c \lambda_{\min}(\Sigma|A) \frac{w(A)}{\sqrt{n}} \leq \inf_{u \in A} \frac{1}{n} \|Xu\|^2 . \tag{142}$$

Putting (141) and (142) together completes the proof. ∎

E Generalized Linear Models: Restricted Strong Convexity

We establish bounds on the regularization parameter and RSC condition for GLMs as discussed in Section 6, along with a few specific examples.

E.1 Generalized Linear Models

Loss functions for GLMs are derived as maximum likelihood estimators for the family of exponential distributions. The canonical density function of exponential family distributions is given by [4, 8, 44]:

$$P(y|\eta) \propto \exp\{\eta y - \varphi(\eta)\} , \tag{143}$$

where η is the natural parameter and has a one-to-one function mapping with the mean parameter $\mu = E[y]$ of the distribution, $\varphi(\eta)$ is the log-partition function which ensures that $P(y|\eta)$ remains a probability distribution. The gradient of the log-partition function is the response function, i.e., $g(\cdot) = \varphi'(\cdot)$, which is monotonic by construction. The inverse of the response function is the so-called link function $h(\cdot) = g^{-1}(\cdot)$.

The mean of the distribution can be obtained from the gradient of the log-partition function at the natural parameter, i.e.,

$$\mu = \varphi^{'}(\eta) = g(\eta) \ . \tag{144}$$

The interested reader can study details and additional properties of exponential families from the existing literature [8, 4, 44]. Examples of distributions from the exponential family include the Gaussian, multinomial, exponential, Dirichlet, Poisson, Gamma, etc.

GLMs are obtained from conditional exponential family distributions by assuming a suitable parametric form of the natural parameter η in terms of X and θ^*, in particular $\eta_i = \langle X_i, \theta^* \rangle$. Then, the conditional distribution is given by

$$P(y_i|X_i, \theta^*) \propto \exp\{\eta_i y_i - \varphi(\eta_i)\} = \exp\{\langle X_i, \theta^* \rangle y_i - \varphi(\langle X_i, \theta^* \rangle)\} \ . \tag{145}$$

The loss function for GLMs simply consider the negative log likelihood of such conditional exponential family forms. Assuming samples to be independent, we have

$$\mathcal{L}(\theta; Z^n) = -\frac{1}{n}\sum_{i=1}^{n}\{\eta_i y_i - \varphi(\eta_i)\} = -\frac{1}{n}\sum_{i=1}^{n}\{\langle X_i y_i, \theta \rangle - \varphi(\langle X_i, \theta \rangle)\} \ . \tag{146}$$

Using chain rule, the first derivative of the loss function evaluated at θ^* is

$$\nabla_\theta \mathcal{L}(\theta^*; Z^n) = -\frac{1}{n}\sum_{i=1}^{n} y_i X_i + \frac{1}{n}\sum_{i=1}^{n} X_i \frac{\partial \varphi(\langle \theta^*, X_i \rangle)}{\partial \eta_i} = \frac{1}{n}\sum_{i=1}^{n} X_i (E[y_i|X_i] - y_i) = \frac{1}{n} X^T \omega \ ,$$

where each element of $\omega \in R^n$ is given as $\omega_i = E(y|X_i) - y_i$. Next we look at some specific examples of exponential families and corresponding GLMs.

1. Gaussian distribution: If the variance of the Gaussian distribution $P(y|X_i)$ is assumed to be 1, then we have

$$P(y_i|X_i; \theta^*) \propto \exp\left\{ y_i \langle X_i, \theta^* \rangle - \frac{\langle X_i, \theta^* \rangle^2}{2} \right\} \ .$$

   Comparing it with the canonical form given earlier, the natural parameter is $\eta_i = \langle X_i, \theta^* \rangle$, log-partition function is $\varphi(\langle X_i, \theta^* \rangle) = \frac{\langle X_i, \theta^* \rangle^2}{2}$ and hence $E[y|X_i] = \varphi^{'}(\langle X_i, \theta^* \rangle) = \langle X_i, \theta^* \rangle$. The noise $\omega_i = y_i - E(y_i|X_i)$ is Gaussian. Considering the negative log-likelihood, the GLM corresponding to the Gaussian distribution yields least squares regression [8].

2. Bernoulli distribution: Assuming the conditional distribution of $y_i|X_i, \theta^*$ to have a Bernoulli distribution with conditional mean parameter p_i, which is a suitable function of $\langle X_i, \theta^* \rangle$, the likelihood of the observations is given by

$$P(y_i|p_i) = p_i^{y_i}(1-p_i)^{(1-y_i)} = \exp(y_i \log p_i + (1-y_i)\log(1-p_i)) = \exp\left(y_i \log\left(\frac{p_i}{1-p_i}\right) + \log(1-p_i) \right)$$

   Therefore, the natural parameter $\eta_i = \langle X_i, \theta^* \rangle = \log\left(\frac{p_i}{1-p_i}\right)$ giving $p_i = \frac{\exp(\langle X_i, \theta^* \rangle)}{1+\exp(\langle X_i, \theta^* \rangle)}$. It can be verified from the fact that the probability density function $P(y_i|p_i)$ adds to 1 and the log-partition function evaluates to $\varphi(\eta_i) = \log(1-p_i) = \log(1+\exp(\langle X_i, \theta^* \rangle))$. The noise in the model corresponds to random draws from a Bernoulli distribution, and each element of ω is $\omega_i = p_i - y_i = \frac{\exp(\langle X_i, \theta^* \rangle)}{1+\exp(\langle X_i, \theta^* \rangle)} - y_i$, which is bounded and hence sub-Gaussian. Considering the negative log-likelihood, the GLM corresponding to the Bernoulli distribution yields logistic regression [8].

3. Poisson distribution: Assuming the conditional distribution of $y_i|X_i, \theta^*$ to have a Poisson distribution with conditional mean parameter λ_i, which is a suitable function of $\langle X_i, \theta^* \rangle$, the likelihood of the observations is given by

$$P(y_i|\lambda_i) = \frac{\lambda_i^{y_i} \exp(-\lambda_i)}{y_i!} \propto \exp\{\log(\lambda_i^{y_i} \exp(-\lambda_i))\} = \exp\{y_i \log \lambda_i - \lambda_i\} ,$$

where the $1/y_i!$ term constitutes the base measure for the distribution. Based on the form, the natural parameter $\eta_i = \langle X_i, \theta^* \rangle = \log \lambda_i$ giving $\lambda_i = \exp(\eta_i) = \exp(\langle X_i, \theta^* \rangle)$. Also it can be verified that the log-partition function $\varphi(\eta_i) = \exp(\eta_i) = \lambda_i$. Each element of ω is $\omega_i = \lambda_i - y_i = \exp(\langle X_i, \theta^* \rangle) - y_i$. Considering the negative log-likelihood, the GLM corresponding to the Poisson distribution yields Poisson regression [8].

As discussed in Section 6, if the design matrix X and the noise is assumed to be sub-Gaussian, then the regularization parameter λ_n needs to be $O(\frac{\Omega_R}{\sqrt{n}})$, following the analysis and results in Section 3. In the rest of this section, we focus on proving the GLMs with sub-Gaussian designs satisfy the RSC condition with sample complexity depending on the width of the spherical cap corresponding to the error set, as discussed in Section 2.

E.2 RSC condition for GLMs

For any convex loss function to satisfy the RSC condition on any $A \subseteq S^{p-1}$, the following inequality

$$\delta\mathcal{L}(\theta^*, u; Z^n) = \mathcal{L}(\theta^* + u; Z^n) - \mathcal{L}(\theta^*; Z^n) - \langle \nabla\mathcal{L}(\theta^*; Z^n), u \rangle \geq \kappa \|u\|_2^2 \tag{147}$$

needs to hold $\forall u \in A_r$. For the general formulation of GLM discussed earlier, we have

$$\begin{aligned}\delta\mathcal{L}(\theta^*, u; Z^n) = &-\langle \theta^* + u, \frac{1}{n}\sum_{i=1}^n y_i X_i \rangle + \frac{1}{n}\sum_{i=1}^n \varphi(\langle \theta^* + u, X_i \rangle) + \langle \theta^*, \frac{1}{n}\sum_{i=1}^n y_i X_i \rangle \\ &- \frac{1}{n}\sum_{i=1}^n \varphi(\langle \theta^*, X_i \rangle) - \langle -\frac{1}{n}\sum_{i=1}^n y_i X_i + \frac{1}{n}\sum_{i=1}^n X_i \varphi^{'}(\langle \theta^*, x_i \rangle), u \rangle .\end{aligned}$$

Simplifying the expression and applying mean value theorem twice we get the following

$$\delta\mathcal{L}(\theta^*, u; Z^n) = \frac{1}{n}\sum_{i=1}^n \varphi^{''}\left(\langle \theta^*, X_i \rangle + \gamma_i \langle u, X_i \rangle\right) \langle u, X_i \rangle^2 , \tag{148}$$

for suitable $\gamma_i \in [0, 1]$. The RSC condition for GLMs then needs to consider lower bounds for

$$\delta\mathcal{L}(\theta^*, u; Z^n) = \frac{1}{n}\sum_{i=1}^n \varphi^{''}(\langle \theta^*, X_i \rangle + \gamma_i \langle u, X_i \rangle) \langle u, X_i \rangle^2 \tag{149}$$

where $\gamma_i \in [0, 1]$. The second derivative of the log-partition function is always positive. Since the RSC condition relies on a non-trivial lower bound for the above quantity, the analysis will suitably consider a compact set where $\ell = \ell_\varphi(T) = \min_{|a| \leq 2T} \varphi^{''}(a)$ is bounded away from zero. The only assumption outside this compact set $\{a : |a| \leq 2T\}$ is that the second derivative is greater than 0. Further, we assume $\|\theta^*\|_2 \leq c_1$ for some constant c_1. With these assumptions

$$\delta\mathcal{L}(\theta^*, u; Z^n) \geq \frac{\ell}{n}\sum_{i=1}^n \langle X_i, u \rangle^2 \mathbb{I}[|\langle X_i, \theta^* \rangle| < T] \mathbb{I}[|\langle X_i, u \rangle| < T] . \tag{150}$$

We give a characterization of the RSC condition for isotropic sub-Gaussian design matrices $X \in \mathbb{R}^{n\times p}$. We consider $u \in A \subseteq S^{p-1}$ so that $\|u\|_2 = 1$. Further, we assume $\|\theta^*\|_2 \le c_1$ for some constant c_1. Assuming X has sub-Gaussian rows with $|||X_i|||_{\varphi_2} \le \kappa$, $\langle X_i, \theta^*\rangle$ and $\langle X_i, u\rangle$ are sub-Gaussian random variables with sub-Gaussian norm at most $C\kappa$.

Let ε_1 and ε_2 denote the tail probability that $\langle X_i, u\rangle$ and $\langle X_i, \theta^*\rangle$ exceeds some constant T, i.e., $\varepsilon_1(T;u) = P\{|\langle X_i, u\rangle| > T\} \le e\cdot\exp(-c_2T^2/C^2\kappa^2) = \bar{\varepsilon}_1$, and $\varepsilon_2(T;\theta^*) = P\{|\langle X_i, \theta^*\rangle| > T\} \le e\cdot\exp(-c_2T^2/C^2\kappa^2) = \bar{\varepsilon}_2$, where $\bar{\varepsilon}_1 = \bar{\varepsilon}_1(T,\kappa)$ and $\bar{\varepsilon}_2 = \bar{\varepsilon}_2(T;\kappa)$ are uniform upper bounds on the individual tail probabilities. The result we present below is in terms of the above defined constants $\ell = \ell_\varphi(T)$, $\bar{\varepsilon}_1 = \bar{\varepsilon}_1(T,\kappa)$ and $\bar{\varepsilon}_2 = \bar{\varepsilon}_2(T,\kappa)$ for any suitably chosen T.

Theorem 7 *Let $X \in \mathbb{R}^{n\times p}$ be a design matrix with independent isotropic sub-Gaussian rows such that $|||X_i|||_{\varphi_2} \le \kappa$. Then, for any set $A \subseteq S^{p-1}$ for suitable constants $\eta, c > 0$, with probability at least $1 - 2\exp\left(-\eta w^2(A)\right)$, we have*

$$\inf_{u\in A} \partial\mathcal{L}(\theta^*; u, X) \ge \ell\underline{\rho}^2\left(1 - c\kappa_1^2\frac{w(A)}{\sqrt{n}}\right) . \tag{151}$$

where $\underline{\rho}^2 = \inf_{u\in A}\rho_u^2$, with $\rho_u^2 = E[\langle X_i, u\rangle^2\mathbb{I}[|\langle X_i, \theta^*\rangle| < T]\mathbb{I}[|\langle X_i, u\rangle| < T]]$, and $\kappa_1 = \frac{\kappa}{(1-\bar{\varepsilon}_1-\bar{\varepsilon}_2)^2}$.

Proof: For any fixed T, let $\bar{Z}_i = \bar{Z}_i^u = \langle X_i, u\rangle\mathbb{I}(|\langle X_i, u\rangle| \le T)\mathbb{I}(|\langle X_i, \theta^*\rangle| \le T)$. Then, the probability distribution over $\bar{Z}_i$ can be written as:[3]

$$P(\bar{Z}_i = z) = \frac{P(\langle X_i, u\rangle = z)\mathbb{I}(|\langle X_i, u\rangle| \le T)\mathbb{I}(|\langle X_i, \theta^*\rangle| \le T)}{P(|\langle X_i, u\rangle| \le T, |\langle X_i, \theta^*\rangle| \le T)} \le \frac{1}{1-\bar{\varepsilon}_1-\bar{\varepsilon}_2}P(\langle X_i, u\rangle = z) . \tag{152}$$

As a result, $\left|\left|\left|\bar{Z}_i\right|\right|\right|_{\psi_2} \le \frac{\kappa}{1-\bar{\varepsilon}_1-\bar{\varepsilon}_2} = \kappa_1$. Thus, $\bar{Z}_i = \bar{Z}_i^u$ is a sub-Gaussian random variable for any $u \in A$. Let $\rho_u^2 = E[(\bar{Z}_i^u)^2] > 0$. Let X_0 be i.i.d. as the rows $X_i, i = 1, \ldots, n$. Let $A \subseteq S^{p-1}$ and consider the following class of functions: $F = \{\frac{1}{\rho_u}\langle\cdot, u\rangle\mathbb{I}(|\langle\cdot, u\rangle| \le T)\mathbb{I}(|\langle\cdot, \theta^*\rangle| \le T) : u \in A\}$. Then for any $f \in F$, $f(X_0) = \frac{1}{\rho_u}\langle X_0, u\rangle\mathbb{I}(|\langle X_0, u\rangle| \le T)\mathbb{I}(|\langle X_0, \theta^*\rangle| \le T)$ and, by construction, F is a subset of the unit sphere, i.e., $F \subseteq S_{L_2}$. Further, $\sup_{f\in F}|||f|||_{\psi_2} \le \kappa_1/2$.

Next, we show that for the current setting, the γ_2-functional can be upper bounded by $w(A)$, the Gaussian width of A. Since the process is sub-Gaussian with φ_2-norm bounded by κ_1, we have

$$\gamma_2(F \cap S_{L_2}, |||\cdot|||_{\psi_2}) \le \kappa_1\gamma_2(F \cap S_{L_2}, |||\cdot|||_{L_2}) \le \kappa_1 c_4 w(A) , \tag{153}$$

where the last inequality follows from generic chaining, in particular [37, Theorem 2.1.1], for an absolute constant $c_4 > 0$.

In the context of Theorem 10, we choose

$$\theta = c_1c_4\kappa_1^2\frac{w(A)}{\sqrt{n}} \ge c_1\kappa_1\frac{\gamma_2(F \cap S_{L_2}, |||\cdot|||_{\varphi_2})}{\sqrt{n}}, \tag{154}$$

so that the condition on θ is satisfied. With this choice of θ, we have

$$\theta^2 n/\kappa_1^4 = c_1^2c_4^2w^2(A) . \tag{155}$$

Then, from Theorem 10, it follows that with probability at least $1 - \exp(-\eta w^2(A))$, we have

$$\sup_{u\in A}\left|\frac{1}{\rho_u n}\sum_{i=1}^n\langle X_i, u\rangle^2\mathbb{I}(|\langle X_0, u\rangle| \le T)\mathbb{I}(|\langle X_0, \theta^*\rangle| \le T) - 1\right| \le c\kappa_1^2\frac{w(A)}{\sqrt{n}} \tag{156}$$

[3]With abuse of notation, we treat the distribution over $\bar{Z}_i$ as discrete for ease of notation. A similar argument applies for the true continuous distribution, but more notation is needed.

where $\eta = c_2 c_1^2 c_4^2$ and $c = c_1 c_2$ are absolute constants. Thus, with probability at least $1 - \exp(-\eta w^2(A))$,

$$\inf_{u \in A} \frac{1}{n} \sum_{i=1}^{n} \langle X_i, u \rangle^2 \mathbb{I}(|\langle X_0, u \rangle| \leq T) \mathbb{I}(|\langle X_0, \theta^* \rangle| \leq T) \geq \inf_{u \in A} \ \rho_u^2 \left(1 - c\kappa_1^2 \frac{w(A)}{\sqrt{n}}\right) . \tag{157}$$

Denoting $\underline{\rho}^2 = \inf_{u \in A} \rho_u^2$, with probability at least $1 - \exp(-\eta w^2(A))$, we have

$$\inf_{u \in A} \partial \mathcal{L}(\theta^*; u, X) \geq \inf_{u \in A} \ \frac{\ell}{n} \sum_{i=1}^{n} \langle X_i, u \rangle^2 \mathbb{I}[|\langle X_i, \theta^* \rangle| < T] \mathbb{I}[|\langle X_i, u \rangle| < T] \geq \ell \underline{\rho}^2 \left(1 - c\kappa_1^2 \frac{w(A)}{\sqrt{n}}\right) . \tag{158}$$

That completes the proof. ∎

Acknowledgements: We thank the reviewers of the conference version [2] for helpful comments and suggestions on related work. We thank Sergey Bobkov, Snigdhansu Chatterjee, and Pradeep Ravikumar for helpful discussions related to the paper. The research was supported by NSF grants IIS-1447566, IIS-1422557, CCF-1451986, CNS-1314560, IIS-0953274, IIS-1029711, and by NASA grant NNX12AQ39A.

References

[1] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. Living on the edge: Phase transitions in convex programs with random data. *Inform. Inference*, 3(3):224–294, 2014.

[2] A. Banerjee, S. Chen, F. Fazayeli, and V. Sivakumar. Estimation with norm regularization. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.

[3] A Banerjee, S Merugu, I Dhillon, and J Ghosh. Clustering with Bregman Divergences. *Journal of Machine Learning Research*, 6:1705–1749, 2005.

[4] O Barndorff-Nielsen. *Information and Exponential Families in Statistical Theory*. John Wiley and Sons, 1978.

[5] P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.

[6] P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.

[7] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

[8] L Brown. *Fundamentals of Statistical Exponential Families*. Institute of Mathematical Statistics, 1986.

[9] P. Buhlmann and S. van de Geer. *Statistics for High Dimensional Data: Methods, Theory and Applications*. Springer Series in Statistics. Springer, 2011.

[10] E. Candes. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathematique*, 346(9-10):589–592, 2008.

[11] E. Candes and T Tao. The Dantzig selector: statistical estimation when p is much larger than n. *The Annals of Statistics*, 35(6):2313–2351, 2007.

[12] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52:489–509, 2006.

[13] E. J. Candes and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51:4203–4215, 2005.

[14] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.

[15] S. Chatterjee, S. Chen, and A. Banerjee. Generalized dantzig selector: Application to the k-support norm. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.

[16] S. Chen and A. Banerjee. Structured estimation with atomic norms: General bounds and applications. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.

[17] Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Struct. Algorithms*, 22(1):60–65, 2003.

[18] M. A. T. Figueiredo and R. D. Nowak. Sparse estimation with strongly correlated variables using ordered weighted l1 regularization. *arXiv:1409.4005*, 2014.

[19] Y. Gordon. Some inequalities for gaussian processes and applications. *Israel Journal of Mathematics*, 50(4):265–289, 1985.

[20] Y. Gordon. On milmans inequality and random subspaces which escape through a mesh in r^n. In *Geometric Aspects of Functional Analysis*, volume 1317 of *Lecture Notes in Mathematics*, pages 84–106. Springer, 1988.

[21] B. Klartag and S. Mendelson. Empirical processes and random projections. *Journal of Functional Analysis*, 225(1):229–245, 2005.

[22] V. Koltchinskii and S. Mendelson. Bounding the smallest singular value of a random matrix without concentration. *arXiv:1312.3580*, 2013.

[23] G. Lecué and S. Mendelson. Sparse recovery under weak moment assumptions. *arXiv:1401.2188*, 2014.

[24] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, 2013.

[25] A. Maurer, M. Pontil, and B. Romera-Paredes. An Inequality with Applications to Structured Sparsity and Multitask Dictionary Learning. In *Conference on Learning Theory (COLT)*, 2014.

[26] N. Meinshausen and B. Yu. Lasso-type recovery of sparse representations for high-dimensional data. *The Annals of Statistics*, 37(1):246—270, 2009.

[27] S. Mendelson. Learning Without Concentration. In *Journal of the ACM*, 2015.

[28] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Reconstruction and subGaussian operators in asymptotic geometric analysis. *Geometric and Functional Analysis*, 17:1248–1282, 2007.

[29] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for the analysis of regularized M-estimators. *Statistical Science*, 27(4):538–557, 2012.

[30] R. I. Oliveira. The lower tail of random quadratic forms, with applications to ordinary least squares and restricted eigenvalue properties. *arXiv:1312.2903*, 2013.

[31] S. Oymak, C. Thrampoulidis, and B. Hassibi. The Squared-Error of Generalized Lasso: A Precise Analysis. *arXiv:1311.0830*, 2013.

[32] Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory*, 59(1):482–494, 2013.

[33] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted Eigenvalue Properties for Correlated Gaussian Designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.

[34] Z. Rudelson and S. Zhou. Reconstruction from anisotropic random measurements. *IEEE Transactions on Information Theory*, 59(6):3434–3447, 2013.

[35] Walter Rudin. *Principles of Mathematical Analysis*. International Series in Pure & Applied Mathematics. McGraw-Hill, 3rd edition, 1976.

[36] V. Sivakumar, A. Banerjee, and P. Ravikumar. Beyond sub-gaussian measurements: High-dimensional structured estimation with sub-exponential designs. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.

[37] M. Talagrand. *The Generic Chaining*. Springer, 2005.

[38] M. Talagrand. *Upper and Lower Bounds for Stochastic Processes*. Springer, 2014.

[39] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.

[40] J. A. Tropp. Convex recovery of a structured signal from independent random linear measurements. In *Sampling Theory, a Renaissance*. (To Appear), 2015.

[41] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y. Eldar and G. Kutyniok, editors, *Compressed Sensing*, chapter 5, pages 210–268. Cambridge University Press, 2012.

[42] R. Vershynin. Estimation in high dimensions: A geometric perspective. *Submitted*, 2014.

[43] M. J. Wainwright. Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ_1-constrained quadratic programming(Lasso). *IEEE Transactions on Information Theory*, 55:2183–2202, 2009.

[44] M J Wainwright and M I Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.

[45] P. Zhao and B. Yu. On model selection consistency of Lasso. *Journal of Machine Learning Research*, 7:2541–2567, November 2006.

[46] S. Zhou. Restricted eigenvalue conditions on subgaussian random matrices. Technical report, Department of Mathematics, ETH Zurich, December 2009.